

Claremont Colleges

Scholarship @ Claremont

CGU Theses & Dissertations

CGU Student Scholarship

Fall 2019

Ethics, X-Phi, and the Expanded Methodological Toolbox: How the Think Aloud Method and Interview Reveal People's Judgments on Issues in Ethics and Beyond

Kyle Thompson
Claremont Graduate University

Follow this and additional works at: https://scholarship.claremont.edu/cgu_etd



Part of the [Philosophy Commons](#)

Recommended Citation

Thompson, Kyle. (2019). *Ethics, X-Phi, and the Expanded Methodological Toolbox: How the Think Aloud Method and Interview Reveal People's Judgments on Issues in Ethics and Beyond*. CGU Theses & Dissertations, 359. https://scholarship.claremont.edu/cgu_etd/359.

This Open Access Dissertation is brought to you for free and open access by the CGU Student Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CGU Theses & Dissertations by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

Ethics, X-Phi,
and the Expanded Methodological Toolbox:

How the Think Aloud Method and Interview
Reveal People's Judgments on Issues in Ethics and Beyond

by

Kyle Thompson

Claremont Graduate University

2019

Approval of the Dissertation Committee

This dissertation has been duly read, reviewed, and critiqued by the Committee listed below, which hereby approves the manuscript of Kyle Thompson as fulfilling the scope and quality requirements for meriting the degree of Doctor of Philosophy in Philosophy.

Dr. Eric Schwitzgebel, Chair
University of California, Riverside
Professor of Philosophy

Dr. Brian Keeley
Pitzer College
Professor of Philosophy

Dr. Phil Zuckerman
Pitzer College
Professor of Sociology and Secular Studies

Abstract

Ethics, X-Phi,
and the Expanded Methodological Toolbox:

How the Think Aloud Method and Interview
Reveal People's Judgments on Issues in Ethics and Beyond

by
Kyle Thompson

Claremont Graduate University: 2019

Ethics isn't a conversation exclusive to philosophers. There is value, then, in not only understanding how laypeople think about issues in ethics, but also bringing their judgments into dialogue with those of philosophers in order to make sense of agreement, disagreement, and the consequences of each. Experimental philosophers facilitate this dialogue uniquely by capturing laypeople's judgments and analyzing them in light of philosophical theory. They have done so almost exclusively by using face valid quantitative surveys about philosophically interesting thought experiments. Based on high participant support for this or that response, researchers conclude that a given theory is more or less intuitive, and in some cases that it is or isn't true. However, such conclusions can only be drawn from quantitative survey data if one assumes those data accurately reflect laypeople's thinking on the issues of interest, an assumption that can't be justified by the quantitative data itself since they are potentially opaque. This is the methodological problem of experimental philosophy: experimental philosophers who use quantitative methods alone can't sufficiently demonstrate their data to be reflective of relevant judgments/intuitions. To better explore people's judgments on issues in ethics *and* to determine whether experimental philosophers are or aren't getting what they take themselves to be getting in surveys, I recreated

part of an experiment by Chituc, Henne, Sinnott-Armstrong, and De Brigard (Chituc et al.) on the Kantian principle that “ought implies can” (OIC), or that whenever an agent ought to do something then it is the case that she can, and added a layer of qualitative data to it by having participants think aloud while completing the surveys and conduct a follow-up interview. With similar quantitative results to those from Chituc et al., I argue that my qualitative results 1) cast doubt on Chituc et al.’s data being reflective of many people’s judgments, especially on the more interesting questions relating to moral obligation; 2) reveal unique insights into how people think about OIC; 3) generally speak to the value of experimental philosophers using triangulation of methods (including qualitative ones) in better understanding laypeople’s complex and varied thinking about questions important to ethics and beyond.

Supplemental documents: Appendices A, B, C, D, and E are surveys I used in my study. Appendices C and D are derived from Chituc et al.’s publication “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can.’” Appendix F is the recruitment flyer I distributed for my study. Appendix H is the quantitative data from my study.

I dedicate this dissertation to the three most influential moral philosophers in my life:

Jessie Lyn Thompson, Don Thompson, and Nancy Thompson.

Acknowledgments

I am so grateful for all of the very helpful and very specific feedback that Dr. Masahiro Yamada, Dr. Brian Keeley, and Dr. Phil Zuckerman provided on my work. Their conversations, emails, and comments helped me formulate my ideas with more clarity and accuracy. As much work as this has been, it would have been far more challenging had they not consistently steered me in good directions.

I am truly indebted to Eric Schwitzgebel for generously agreeing to chair my committee after the Philosophy Department at CGU tragically closed. Not only was his guidance extremely helpful, but he also made sure to treat me as a graduate student just like the others he worked with. He made me feel included just when my circumstances made me feel otherwise.

I would also like to thank Geoffrey Weiss for providing me with nothing but the most honest, harsh, and cutting criticism which, despite the intentions behind such remarks, make me a better thinker, writer, and philosopher. He was even useful in helping me chop some logic. And thanks to Jordan Neidlinger for helpful clarification on the trickier aspects of logic. Also, to Marin Laukka for pointing me to some excellent resources and for translating my methodological worries into the appropriate terms of art in contemporary psychology. And I want to thank Ryan Falcioni for chatting with me about my ideas and philosophy more generally; I always came away with some good insights based on our conversations. Also, I'd like to thank Jonathan Weinberg for a helpful conversation about my dissertation and x-phi more generally.

Many thanks to James Griffith for patiently answering my many research questions and guiding me during my first experimental study. Your advice on the theoretical and practical was indispensable.

I am grateful also to the x-phi community for the interesting and insightful work they produce. In particular I want to express the deepest gratitude toward Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard for creating the study that motivated my project. While I criticize aspects of their work, I do so because I believe their achievements are interesting and worthwhile; I do so out of a respectful, mutual quest for truth. And in that pursuit, I want to also thank all the wonderful participants who helped make this study what it is.

Mom. Dad. Thank you. Your support has meant the world to me. None of this would be possible without your love, encouragement, and emphasis on education. Thank you for everything.

I am also very thankful to my wife, Jessie Lyn Thompson, for allowing me to constantly think aloud—yes, in just the same sense as my research!—as I struggled to mold my ideas into something coherent and significant. Her responses to my verbalizations consistently challenged me to be more clear and detailed in constructing and completing my project.

And lastly, to my beloved animals: Koda, Keira, Freya, and Fiona. They helped remind me that at the end of the day, I'm just another animal that needs to play, eat, sleep, and be social.

Table of Contents

CHAPTER ONE

Ethics and X-Phi: An Overview of My Project.....1

CHAPTER TWO

The Methodological Problem of X-Phi:
A Case Study in Free Will and Moral Responsibility.....15

CHAPTER THREE

A Fine-Grained Look at the Methodological Problem:
The “Ought Implies Can” Study Motivating My Work.....61

CHAPTER FOUR

Why I Crafted My “Ought Implies Can” Study the Way
I Did: All Roads Lead to Triangulation of Methods.....98

CHAPTER FIVE

Design and Quantitative Results of My Study:
Adding the Think Aloud Method and Interview
to “Ought Implies Can”.....132

CHAPTER SIX

A First Analysis of Qualitative Data:
Quantitative X-Phi Surveys Fail to Capture
Unambiguous, Relevant Judgments.....163

CHAPTER SEVEN

A Second Analysis of Qualitative Data: The Think Aloud
Method and Interview Suggest that OIC Judgments Are
Sensitive to Concreteness, Definitional Distinctions, and More.....202

CHAPTER EIGHT

Conclusion: Ethics is Messy,
Triangulation is Worth the Trouble.....271

APPENDICES.....284

ENDNOTES.....300

REFERENCES.....326

CHAPTER ONE

Ethics and X-Phi:

An Overview of My Project

I. Ethics Meets X-Phi

Ethics is the concern of humanity, not merely philosophers or other experts. We go about our days treating others, other animals, landscapes, and institutions in ways that accord with our consciences, what we deem as good and fair, etc. We theorize and act upon our judgments about what is right, when it is right, and we get bound up in our feelings when we succeed and fail. There is great value, then, in not only learning how laypeople think about issues in ethics, but also bringing their judgments, intuitions, and deliberations into dialogue with those of philosophers. Is there agreement? Disagreement? What is the significance of each? Ought philosophers maintain ethical theories that don't accord with laypeople's understanding of common concepts like the moral responsibility bestowed by free will, and the nature of moral obligation? What does it mean to learn that pervasive judgments about this or that thought experiment are shared with one cultural community but not another? The answer to these is in experimental philosophy (x-phi).

Experimental philosophers take these questions seriously by running studies to capture the ways laypeople think about issues in ethics before bringing their findings into a more traditional conversation in ethics. The most persuasive argument for such research, ironically, comes from what is perhaps the most common complaint against it. The complaint is that experts in other fields don't worry themselves with the judgments/intuitions of laypeople in solving problems at the hearts of their disciplines, so there is little good in philosophers doing so. The rebuttal is that some questions in philosophy, though handled by experts, are straightforwardly borne out of and

have their cash value in ordinary human living, not in a remote lab or in a philosophy so esoteric that only specialists who've trained for years can speak to them. Joshua Knobe and Shaun Nichols, in "An Experimental Philosophy Manifesto," best articulate this kind of argument. They concede that this complaint applies to some of philosophy, for example "if we want to know whether the representational theory of mind is superior to connectionist alternatives" then it wouldn't serve us to run studies on laypeople's judgments, but for other philosophy, such as "problems concerning free will, personal identity, knowledge, and morality" it matters what laypeople think—in fact, common judgments and intuitions define such philosophical questions.¹ In this project, then, I look at and participate in experimental philosophy (x-phi) to better understand moral responsibility (as bound by free will) and moral obligation. These are debates that aren't at all alien to common philosophizing, and therefore it is important to learn how laypeople think about them while crafting theories in their service.

The problem, though, is that the current practices of x-phi give me pause. Surely, many psychologists and philosophers alike have read the claims of a research publication in psychology and objected to the evidentiary support that was attributed to some survey, some experiment, or some test. In particular, many have looked at statistically significant experimental results derived from diverse populations and thought something like, "I just don't see how *these* survey questions measure this complex phenomenon even if the data are reliably produced." It is from this general picture that my project emerges: I found myself reading studies in experimental philosophy (x-phi)—in which philosophical ideas are studied via the methods of the social sciences—and thinking out loud, "I just don't see how these questions tap a person's judgment regarding this complex philosophical issue, successful replication aside." This was especially the case with

questions in ethics, which manifest in everyday conversations with richness, complexity, and nuance.

But such an utterance from the armchair does little to advance scholarly discussion. So, my goal in this project has been simple: run a study that explores some psychological phenomenon (in this case a philosophical judgment about ethics), the data used to represent it (in this case one or more quantitative values selected in response to one or more survey questions), and the relationships between the phenomenon and the data. To do so, I employed two qualitative methods—the think aloud method² and the interview—to learn more about the potentially messy cognitive processes that occur when someone reads and completes a survey for x-phi research in ethics, the cognitive processes that should “be” or “contain” the relevant philosophical judgments of interest. The first method has subjects talking aloud as they complete a traditional x-phi survey, with the hope that their utterances reveal their thinking as they complete the survey, and the second has subjects answering interview questions about the survey they completed, with the hope that this too will be revelatory or at least clarificatory of their responses to the survey.

Much in the way of experimental philosophy—which metaphorically sets fire to the armchair³—I stood up, ran a study, and now I present this project as an in-depth look at my reasons for doing so and what I found. I chose to focus on the philosophical assertion that “ought implies can” (OIC)—or the notion (connected to Kant⁴) that moral obligation only obtains for actions that agents *can* carry out—but I could have looked just the same at other issues connected to ethics and beyond, such as those in free will and moral responsibility. (In fact, I spend chapter two looking at the experimental philosophy of free will and moral responsibility before moving exclusively to OIC, in part to show that my project speaks to multiple philosophical debates.) And, in sum, I argue that my armchair skepticism was *not* unfounded. Specifically, as my core thesis, I argue

that 1) much of x-phi suffers from a methodological problem in that the extensive use of quantitative surveys to capture people's philosophical judgments cannot itself justify belief that such surveys are capturing relevant judgments; 2) employing qualitative methods in x-phi research, namely the think aloud method and the interview, reveals unique insights into people's philosophical judgments about issues in ethics and thus insights into the debates surrounding the issues concerned with these judgments; 3) employing these methods also exposes unique pitfalls of x-phi studies that rely solely or predominantly on quantitative survey methods; 4) all of the previous sub-claims suggest that triangulation of methods is worth pursuing both for developing x-phi surveys and for better capturing judgments related to ethics and beyond, which are more rich and nuanced than often assumed.

I hope this project not only enlivens debate in both philosophy and psychology, but that it also inspires more of us to see the value in placing an armchair adjacent to the laboratory, if it hasn't already been turned to ashes. When it comes to moral philosophy, humanity has much to offer that is thoughtful, nuanced, and interesting—all we have to do is ask them to share it.

II. My Thesis at the Chapter Level

Points 1-4 listed above constitute the core thesis of this project. However, these points do not map onto the structure of my chapters, nor do they indicate some of the more fine-grained details of this project that would help a reader navigate from chapter to chapter. What follows is just that.

- CHAPTER ONE: Rounding out this current chapter is a list of definitions of key terminology and abbreviations.

- CHAPTER TWO: I introduce the main problem I see with much x-phi research that employs thought experiments to capture subjects' philosophical judgments/intuitions. This problem, which I term "the methodological problem of experimental philosophy," is that experimental philosophers (or x-philosophers) generally do not have sufficient evidence to show that quantitative survey responses reflect relevant philosophical judgments. I focus on experiments in the x-phi of free will and moral responsibility in this chapter and then spend the rest of the dissertation looking at experiments on "ought implies can"; the goal is to show that my methodological concerns aren't specific to any topic in experimental philosophy.
- CHAPTER THREE: I take a fine-grained look at the methodological problem of experimental philosophy as it is apparent in the study on "ought implies can" that I go on to partially re-create. This gives me the opportunity to highlight the ideal story of data collection (i.e., what researchers using surveys ideally expect to happen in their studies) as well as introduce the details of the study I based my study on.
- CHAPTER FOUR: I discuss why I crafted my own study the way I did, and I explore how it fits into the larger x-phi research program. In a sense, I justify my use of two qualitative methodologies (the think aloud method and the qualitative interview) to learn more about both the data collection process of x-phi studies and the philosophical judgments people have in response to philosophical questions.

- CHAPTER FIVE: I detail the design of my own study on “ought implies can” and then analyze my quantitative results. In brief, I obtained quantitative data similar to those from the study I partially re-created.
- CHAPTER SIX: This chapter represents my negative project in that I cast doubt on using quantitative surveys in isolation. I analyze the qualitative data from my study on “ought implies can” to argue that viewing quantitative results in isolation can lead to misguided conclusions.
- CHAPTER SEVEN: This chapter represents my positive project in that I contribute claims as to how we might better understand “ought implies can.” I use my qualitative data to make seven suggestions about how we might better understand people’s judgments regarding “ought implies can.” The total picture formed from these suggestions is that laypeople’s judgments are complex and sensitive to a range of influences, many borne out in the literature on moral obligation, which points to the value of using triangulation of methods to more fully understand people’s thinking.
- CHAPTER EIGHT: I discuss the limitations of my own work and the future of x-phi methodological practice, and I reiterate my suggestion that experimental philosophy ought use triangulation of methods.

III. Key Terms and Abbreviations Defined

Before proceeding, I will define key terms that I use extensively—such as “laypeople,” “judgments,” “vignette,” “follow-up question,” and “thought experiment”—since not all x-phi literature, or relevant psychology literature, uses identical terminology. After this, I’ve included

some of the abbreviations that I end up using extensively—such as OIC, x-phi, MPEP, the OIC study, and OIC experiment 3.

Laypeople: In brief, I will use the term “layperson,” and the plural “laypeople,” to refer to someone who is not familiar with the terms, debates, theories, and concepts employed in academic discussions of Western analytic philosophy in general or the specific debate around “ought implies can” (OIC) in particular (as this specific debate, which I detail later, is central to my research experiment).⁵

Judgments (used interchangeably with “philosophical judgments”): In brief, I understand a philosophical judgment to be a person’s cognitive response to—or thinking in regard to—a certain philosophical idea, concept, proposition, or theory. This cognitive response can be roughly described as how that person regards the idea, concept, proposition, or theory, or how they evaluate it. For example, my philosophical judgment is that some actions are morally preferable to others because I regard them to be so. As another example, some people, after evaluating different hypothetical scenarios, philosophically judge that people do exercise free will in at least some situations. A judgment, as I understand it, can be formed on the basis of direct engagement with a philosophical idea, concept, proposition, or theory, or on the basis of engagement with any content or scenario that points to or embodies a philosophical idea, concept, proposition, or theory. For example, in directly discussing doing and allowing harm, a person can form the philosophical judgment that doing harm is morally distinct from allowing harm; or, in reading a story that implicitly or indirectly introduces the distinction between doing and allowing harm, a person can form a similar judgment.

I will use “judgments” over “intuitions”⁶ for two main reasons. First, as far as I know, the former isn’t embroiled in the same level of controversy as is the latter within the context of x-phi.

Within the context of x-phi, there is considerable disagreement as to what an intuition is⁷ such that invoking the term within a conversation in x-phi can bring about a wealth of confusion. In my reading of x-phi publications, I can't recall any such controversy around judgments; this makes "judgments" preferable for my project. Second, and relatedly, my reading of publications in and related to x-phi has given me the impression that the term "judgment" is less restrictive in meaning than "intuition," which is important because my project here is not only concerned with people's cognition that can be restrictively understood as intuitions. I am interested, in quite general terms, with how people think about certain philosophical ideas, concepts, propositions, etc. In other words, I prefer "judgment" because I think it comfortably includes more restrictive cognition, such as that which philosophers label "intuitions," without excluding more general cognition that might be excluded on certain definitions of "intuitions." To ground this distinction between "judgment" and "intuition" a bit more, I do want to note the similarity in the concepts and present *some* kind of definition of "intuition" so it is clear to the reader that I'm using "judgment" to mean something more general than an "intuition." I loosely conceive of an "intuition" (or "philosophical intuition") as a spontaneous cognitive response to a philosophical idea, concept, thought experiment, or theory that manifests in a person's cognition as a "seeming that" a particular idea, conclusion, or judgment is one thing or another (e.g., that a judgment is right, or that a hypothetical person in a thought experiment is not acting on her own free will).⁸ But again, I'm opting not to use this term most of the time since other philosophers might reject this definition for one of the many others given that there is widespread disagreement as to what an intuition is. Instead, I will stick to "judgments" which include what I have just defined as an "intuition" along with other cognitive responses, such as those that aren't quite so spontaneous.

Vignette: In my study, a “vignette” refers to a brief, text-based description of some scenario, typically one that is hypothetical. In other contexts, vignettes are more varied—from text to audio to live performance⁹—but here I have the specific text version in mind. Importantly, in my usage, a vignette does not include any questions for consideration (see next term description for more information); it merely describes a scenario. For example, here is a brief vignette I’ve invented that relates to the philosophy of free will: *Amy has observed that Chris purchases a coffee from the vending machine at the office before 9 a.m. every work morning. On Wednesday evening, Amy correctly predicts that Chris will purchase a coffee on Thursday morning before 9 a.m.* Vignettes of this kind are often presented along with one or more follow-up questions that relate to content in the vignette itself.

Follow-up Question: In my research, I understand a follow-up question to be a brief text-based question that is presented in conjunction with a vignette that addresses content in the vignette. Based on the vignette presented in the previous term description, a follow-up question might be: *Does Chris freely choose to buy coffee on Thursday morning?*

Thought Experiment: I understand a thought experiment to be a description—presented textually, visually, orally, etc.—of a hypothetical or actual scenario that is presented along with one or more questions relating to the content of the description. The description and the question are typically designed to cause someone to consider or reconsider an important philosophical idea. For example, the vignette and follow-up question presented in the previous two term descriptions form a thought experiment together.

The OIC study: This refers to the “ought implies can” (OIC) study by Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard, described in their paper entitled

“Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” which details three related experiments.

OIC experiment 3: This refers to the third of three experiments detailed in the OIC study (described above); this is the one I partially recreated in my study that grounds this dissertation.

Chituc et al.: The researchers Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard as connected to the OIC study (described above).

OIC: “ought implies can”; often connected to Kant,¹⁰ the assertion that moral obligation only obtains for those actions that an agent can actually carry out. OK, and that’s the short of it, but there’s a far more complicated debate surrounding it, and a more specific position I need to take as a result (which I will do just below). Of relevance to this project, there is debate as to what relationship truly holds between concepts “ought” and “can” since “implies” is in itself an opaque word. For the purposes of unpacking the OIC debate, I will defer 1) to Walter Sinnott-Armstrong’s mapping out of the OIC principle since his earlier work not only gives rise to the thought experiments used in the study that I partially re-create, but also because he is an author on that very study; and 2) to the specific conversion of this Kantian principle to an empirically testable claim that the researchers of the study I re-create (one of whom is Sinnott-Armstrong) outline in their publications. Then I will articulate how I will be using OIC throughout this project.

First, Sinnott-Armstrong entertains three ways of understanding the relationship between “ought” and “can”: one of entailment (“ought entails can”), one of presupposition (“ought presupposes can”), and one of conversational implicature (“ought conversationally implies can”). Importantly, on entailment, it is false that an agent ought to do something if she cannot do that thing—i.e., the contraposition is in play such that if agent cannot, then it is false that she ought.¹¹ On presupposition, if an agent cannot do something, then it is “*neither true or false*” that she ought

do it¹²—i.e., the claim that she ought *presupposed* that she could, so if she can't then it renders the claim inert in some sense.¹³ On conversational implicature, if an agent cannot do something, Sinnott-Armstrong argues it still could be true that she ought to, depending on features of the situation.¹⁴ This kind of implicature deals with, in so many words, what a speaker and a listener understand is implied in an utterance. Sinnott-Armstrong summarizes Grice: “roughly, saying *p* conversationally implies *q* when saying *p* for a certain purpose cannot be explained except by supposing that the speaker thinks that *q* and thinks that the hearer can figure out that the speaker thinks that *q*, etc.”¹⁵ The example he subsequently gives is that when directing someone to a particular gas station, the speaker’s utterance, which contains directions *and not the hours of operation*, implies to the listener that the speaker thinks the gas station is open *and* that the listener can ascertain from the utterance that the speaker thinks this. And so, if ought conversationally implies can, as Sinnott-Armstrong argues is sometimes the case, then there are cases in which it can be true to say an agent ought to do something she cannot, such as when someone blames an agent for failing to make good on a promised appointment when the appointment comes to pass and while it is clear the agent cannot fulfil it.¹⁶ A quick example that Chituc et al. give of such is the utterance ““Where are you? You ought to be here by now!”” which conveys blame but does not imply “can.”¹⁷

Second, and more important to my project, is how the researchers Chituc et al. think about OIC and convert it to something like an empirically testable principle. In their publication “An Empirical Refutation of ‘Ought’ Implies ‘Can,’” those researchers (in this case Henne et al.) essentially put the entailment version outlined above on the table for testing. (The same researchers put the analytic/conceptual entailment version on the table as well in the other, similar publication, the one I label above as “the OIC study,” but I’ve opted to unpack the Henne et al. publication’s

description as it is clearer.¹⁸) Specifically, they discuss whether the relationship between “ought” and “can” is one of “analyticity” or “conceptual entailment,” an irrelevant distinction for our current purposes, which is essentially to say that “ought” entails “can” just as “bachelor” entails “male”: just as an agent not being male makes it false that she is a bachelor, so too does an agent being unable to do something make it false that she ought to do so.¹⁹ That is, they aren’t directly addressing the presupposition or the conversational implicature views. To convert this OIC principle to an empirical claim, Henne et al. make the following move: “If one claim analytically or conceptually entails another, then competent speakers in good epistemic positions should consistently deny the first when they know that the second is false.”²⁰ That is, if participants in a study consistently deny an agent ought when agent cannot, then this is *consistent with* or *compatible with* OIC.²¹ But if they are happy to ascribe obligation to unable agents, this suggests “ought” doesn’t entail “can.”

So here is *my* stance throughout this dissertation. Unless otherwise stated, I am employing OIC and using language about things that support it, endorse it, etc. to mean the following: that which is at the very least *consistent with* or *compatible with* but perhaps even *positively supportive* of the entailment view and the presupposition view but not *necessarily* a direct and total affirmation of one of those two views such that it excludes the rest (since statements consistent with entailment are also consistent with presupposition and can be with conversational implicature). For example, if I write, ‘This view affirms/supports/endorse/aligns with OIC,’ I am stating at the very least that the view is perfectly compatible with OIC as entailment and presupposition, though it might turn out to be an exclusionary support for one such view in some cases, *and*, depending on the context and verb I employ, I am potentially treating the view as positive support for the truth of OIC. And when I use language of things being against or negating OIC, it is that they are *inconsistent* or

incompatible and therefore *counter* to the entailment view, but not necessarily a rejection of presupposition, conversational implicature, or other. In brief, OIC as entailment is what is on the table here. But the tricky part is that while some utterances directly negate OIC as entailment—e.g., those that affirm obligation for unable agents—many could support *it*, presupposition, or conversational implicature. And I take this stance because 1) it is what Chituc et al. put on the table in their research; and 2) I don't think it is clear enough in laypeople's utterances during my study or in the quantitative questions they are asked in other studies if they are at times endorsing entailment over presupposition or vice versa *even though* it is often clear when they negate OIC or say something that is consistent with entailment, presupposition, or both. One researcher, Mizrahi, argues that his data show that presupposition in particular isn't supported by participant responses after stating that entailment is viewed as too strong a relationship, but it is *also* the case then that his data don't support entailment.²² I don't want to pretend I have fine-grained enough data to locate support for one particular OIC view (when that which is consistent with entailment is consistent with presupposition and can be with conversational implicature), but Mizrahi (and others for that matter) are entitled to take participant responses as being *inconsistent* with entailment and presupposition because, like many propositions, OIC is easier noticed in its violation—e.g., 'He can't at that very moment, but he ought to'—than in its theory-opaque affirmation—e.g., 'No, I don't think he ought to because he can't.'

x-phi: experimental philosophy²³; the recent movement in analytic philosophy that involves employing experimental methodologies of the social sciences to tackle questions associated with philosophy.

MPEP: the methodological problem of experimental philosophy; my invented term which refers to the lack of epistemic justification experimental philosophers have in thinking that a given

quantitative survey in experimental philosophy research accurately elicits and captures a relevant philosophical judgment in subjects.

Nahmias et al.: The researchers Eddy Nahmias, Stephen G. Morris, Thomas Nadelhoffer, and Jason Turner, used in reference to their chapter “Is Incompatibilism Intuitive?” and occasionally their publications “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility” and “The Phenomenology of Free Will.”

The ideal story of data collection in x-phi (ideal story): A term of my creation that refers to what ideally happens in x-phi studies that employ quantitative surveys as participants read and complete a given survey. In brief, the ideal story is that participants form relevant philosophical judgments that are accurately captured, without adulteration, by the follow-up question(s) on a given survey.

The x-phi of OIC: This just refers to the smaller academic debate surrounding “ought implies can” insofar as it is linked with experimental philosophy and empirical research more generally; it is connected to or a sub-debate of the larger philosophical debate on OIC.

Experiment v study: In general, I try to refer to a “study” as an entire experimental project that contains one or more discrete “experiments.”

CHAPTER TWO

The Methodological Problem of X-Phi:

A Case Study in Free Will and Moral Responsibility

I. In X-Phi, Quantitative is King and That's a (Methodological) Problem

You find yourself reading paragraphs written in terse, sterile language: short vignettes about hypothetical characters named Jones and Smith. In the first, you learn about their exceedingly mundane behaviors: they see what appears to be a barn and they claim that it is in fact a barn.²⁴ In the next, you read of their split-second decisions in impossibly rare, high-pressure situations: they choose either to send a train careening into a single innocent person or allow it to collide with five other innocents. Then come the follow-up questions: Did Jones *know* it was a barn? Was it *permissible* for Smith to let the train continue its deadly course?²⁵ You make your selections. While you might not know whether or not you chose the “right” answer—whatever that means—you can rest assured that you were participating in a special type of psychology experiment. You helped create data for research in *experimental philosophy*.

Experimental philosophy, or x-phi²⁶, is a new branch of philosophy that employs the methods of the social cognitive sciences to investigate ideas, theories, and claims traditionally associated with academic philosophy.²⁷ Though experimental philosophers have diverse research goals, they often work to capture and analyze the philosophical judgments or intuitions²⁸ of laypeople²⁹ and philosophers alike in order to better understand the judgments/intuitions themselves as well as answer a variety of questions about the connection between judgments/intuitions and specific philosophical theories, particularly those that appeal to judgments/intuitions for theoretical support.³⁰ For example, some experimental philosophers want

to understand whether commonly accepted philosophical judgments/intuitions are consistently shared by laypeople within or across cultures.³¹ If certain judgments/intuitions are culturally or universally shared, then experimental philosophers can contend that we have more reason to accept theories that reflect these judgments/intuitions. However, if people report a diversity of judgmental/intuitive responses to philosophical ideas, then any theory that relies on a culturally specific judgment/intuition can be called into question as idiosyncratic to a particular culture and thus not universally meaningful. Other experimental philosophers are interested in how laypeople think about important philosophical issues that bear on everyday life, such as free will.³² What connects the diverse projects in x-phi is that experimental philosophers use scientific data that they themselves typically collect to address philosophical issues that rely in one way or another on certain empirical facts about reality or about how people come to think about reality.

And so, experimental philosophers are currently out in the world working on projects that align with the above summary of the x-phi movement—and a good number of them are largely focused on capturing people’s judgments/intuitions for the purpose of informing conversations in philosophy. This is all well and fine, but I worry that the foundation of much x-phi research—i.e., the use of quantitative surveys to capture philosophically interesting judgments/intuitions—might not be methodologically robust enough for experimental philosophers to make the claims they want to make on its behalf. And I am particularly interested in this worry as it applies to ethics, which I think gives rise to some of the most basic yet rich philosophical judgments. But even if this foundation is secure, what is certainly missing from the x-phi movement is a clear sense as to whether experimental philosophers have any real certainty that they are capturing what they intend to capture. I call this worry “the methodological problem of experimental philosophy” (MPEP),

or “the methodological problem” for short. It is the core motivation for this current project insofar as it applies to people’s judgments about issues in ethics.

I approach the x-phi movement in agreement with many of its fundamental tenets; I hold that judgments/intuitions are real, that judgments/intuitions matter for philosophy, and that capturing people’s judgments matters for doing philosophy. However, I argue that what is missing from the movement is sufficient reason to think that many x-phi studies are actually or accurately capturing people’s relevant and coherent judgments. Or, in even briefer terms, I argue that experimental philosophers ought have a clearer sense of what cognitive responses their data reflect. If researchers knew with greater certainty that quantitative surveys *were* successfully capturing people’s relevant judgments, then it would lend greater support to claims made on their behalf, whether prescriptive or descriptive in nature. On the other hand, if these methods were shown to be unreliable or ineffective in capturing said judgments, then it would not only undermine these same claims, but it would also cast doubt on appeals to judgments/intuitions in general. If researchers cannot be sure that they have another person’s relevant judgment in hand, it is difficult to see how philosophers can be sure with their own. This is especially problematic in moral philosophy where strong judgments drive theory and application.

II. The MPEP in the X-Phi of Free Will and Moral Responsibility

In this section, I want to elaborate on the methodological problem of experimental philosophy (MPEP) by pointing to evidence of its existence, ubiquity, and gravity. To do so, I will analyze a case study in the experimental philosophy of free will and moral responsibility. Specifically, I’ll explore a series of related studies discussed in three chapters of *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols. These chapters are ideal for exploring the

MPEP because they capture a self-contained, self-referential experimental analysis and discussion of a topic in philosophy, that of free will and moral responsibility. They also aren't part of the literature on "ought implies can," which will be the focal point of the remainder of my project. This allows me to suggest that the MPEP, arising from methodological practices, affects many topics investigated primarily through quantitative surveys. In analyzing x-phi research in free will/moral responsibility and then later "ought implies can," I will argue that the MPEP broadly and profoundly affects x-phi research that relies heavily on quantitative surveys. My focus is clearly on the x-phi of ethics, given my study on OIC, but in looking at free will, which has both a moral component and a strictly metaphysical one, I hint at the larger implications of my work.

Before cracking open these chapters on the x-phi of free will/moral responsibility, I want to frame my analysis by outlining some important features of the MPEP. At the heart of this methodological problem is, unsurprisingly, a problem of research methodology: because many x-phi studies rely on quantitative surveys alone (or in large part) to capture subjects' *relevant* philosophical judgments (i.e., those judgments that speak to a certain philosophical issue or question that researchers are concerned with), it is unclear whether these studies can be said to capture the specific judgments of interest. Ultimately, this problem is one in which researchers lack evidentiary recourse to justify one of x-phi's most common methodologies: the judgment-capturing quantitative survey. In their search for relevant judgments, experimental philosophers are no doubt actually conducting various experiments and collecting data from subjects. However, it is unclear whether their resultant data accurately reflect people's judgments regarding certain philosophical ideas.

In reality, the methodological problem is a complex problem with potentially many layers. For this current section, I want to note two. First, it is unclear if administering quantitative surveys

about philosophical thought experiments succeeds in reliably *eliciting* the desired cognitive responses in the minds of experimental subjects, namely judgments that are directly related to the philosophical idea(s) being studied. Second, if and when these judgments are successfully elicited, it is unclear if the quantitative surveys often used to capture them do so accurately. In other words, if experimental philosophers want to use their data as evidentiary support for claims about laypeople's judgments (or philosophers' judgments), then they ought to be reasonably sure that when a subject selects the response on a survey that indicates "Jones *knows* that it is a barn" that the subject did in fact judge that Jones has such knowledge *and* that such judging was based on an appropriate understanding, reading, or appraisal of Jones' situation *and* that this is what is captured in the survey response. To recap: as an experimental philosopher, I want a quantitative datum to reflect a survey selection that was caused by a subject's actual judgment based on an actual understanding of the thought experiment I present her with *rather* than being caused by something unlike a judgment or being caused by an actual judgment that results from her non-understanding of the thought experiment. And, importantly, I want the judgment based on that accurate understanding of the thought experiment to be a *relevant* judgment that is, in some sense, *in response* to the features of the thought experiment that I and other researchers are concerned with.

But this is not what researchers actually have in a datum. In many instances, they have a circled response, a clicked button on a computer, etc. that corresponds to "Jones *knows* that it is a barn" and—if I dare speculate—their own background assumptions about what the thought experiment conveys and the ways in which subjects might interpret it. They have this and little more. 'Surely the vignette is *as clear as day*,' some researchers might suppose, 'so what else could be going on?' Maybe it isn't so clear. Perhaps the wording is confusing to laypeople, or it obscures the main idea being explored in the vignette. Or perhaps subjects aren't reacting to the

philosophically interesting features of the survey, or perhaps subjects select the response that they *think* communicates to the researchers that they understand what is being asked of them. ‘In addition,’ some researchers might offer, ‘the subjects succeeded on an attention check prior to completing the survey; they *were* paying attention.’ Yes, but to what? Are we sure they attended to the relevant features of Jones’ situation? Again, what is missing is knowledge *that* experimental philosophers are capturing people’s relevant judgments. To invoke the language of prominent epistemological theory, experimental philosophers may have true belief (i.e., that their data are reflective of people’s judgments), but they lack justification for this belief (i.e., the kind of justification that grants one knowledge in a straightforward sense).³³

Now I’ve only shown that experimental philosophers *may* lack this justification—and thus far my injection of skepticism teeters on the purely speculative. To avoid this, I’ll ground this MPEP analysis in the case study on free will and moral responsibility that I mentioned at the outset of this section. In *Experimental Philosophy*, Knobe and Nichols compile three chapters on the xphi of free will and moral responsibility into a single section entitled “Responsibility, Determinism, and Lay Intuitions”—my focus will be on the middle chapter, by Eddy Nahmias, Stephen G. Morris, Thomas Nadelhoffer, and Jason Turner (hereafter Nahmias et al.), though I will reference the other two extensively. My focus is with this chapter, entitled “Is Incompatibilism Intuitive?”, because it dives into thick philosophical weeds (relating to issues in ethics) in an interesting and useful way for the purposes of outlining the MPEP *and* because the third chapter, by Nichols and Knobe themselves and titled “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions,” directly responds to parts of it in important ways.

Nahmias et al. have a total project that falls squarely within the debate over the problem of free will, and it is informed by x-phi data. They begin their chapter by defining, in an uncontroversial manner, what incompatibilism entails, and by stating what they take themselves to accomplish in the chapter:

Incompatibilists believe that the freedom associated with moral responsibility is impossible if determinism is true, and they often claim that this is the natural view to take given that it is purportedly supported by ordinary intuitions. In this essay, we challenge the claim that incompatibilism is intuitive to most laypersons, and we discuss the significance of this challenge to the free will debate.³⁴

This challenge comes in the form of findings from a series of x-phi-style studies they conduct.³⁵ Now we have a glimpse of Nahmias et al.'s total project and we can appreciate how and to what end they interpret their data: the intuitiveness of incompatibilism (as understood in conjunction with moral responsibility) was tested via the survey methods of x-phi.

I argue that the MPEP in its most basic form is present in the first experiment outlined by Nahmias et al. This form of the MPEP is that of a quantitative survey study whose results plausibly don't reflect a genuine eliciting and capturing of relevant judgments; in this case, subjects' judgments as to whether a hypothetical character, Jeremy, acts freely in a complex situation:

Imagine that in the next century we discover all the laws of nature, and we build a supercomputer which can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any future time. It can look at everything about the way the world is and predict everything about how it will be with 100% accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25, 2150 A.D., twenty years before Jeremy

Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 PM on January 26, 2195. As always, the supercomputer's prediction is correct; Jeremy robs Fidelity Bank at 6:00 PM on January 26, 2195.^{36, 37}

Imagining this were real, participants were asked whether or not Jeremy “acts on his own free will” when he robs the bank; 76% said he did.³⁸ To ensure that it wasn't the immorality of Jeremy's behavior that prompted people to assign freedom to Jeremy's actions, the researchers varied the language for other participants, having Jeremy save a child (still 68% said the act was freely willed) or go jogging (still 79% said the act was freely willed).³⁹

It is true that some might affirm that this vignette generally succeeds in articulating the essence of determinism (as it relates to moral responsibility), leading them to believe the results show participants attributing free will in spite of determinism. However, others might worry that it doesn't succeed, or still others might think that it does but that it also drags with it a number of other potential judgment-nudging oddities. Nothing about this vignette is so clear so as to prevent me from reasonably arguing that subjects might be reacting to a host of philosophically irrelevant or misguided factors. As a point of contrast, I simply wouldn't say the same about a thought experiment such as the following: “Paul murders an innocent person who is walking home from work just because he feels like it; did Paul do something immoral?” That is, I don't take myself to be picking nits in the supercomputer example; certain philosophically interesting questions, such as the Paul example as it serves to pump judgments about the immorality of murder, are straightforward enough that it really isn't epistemologically worrying to conclude from the overwhelming data⁴⁰ in support of Paul behaving immorally that people judge that murdering an innocent stranger is an immoral act. Unlike this Paul example, though, the supercomputer one

warrants skepticism if we think we have straightforward access to people’s judgments based on their quantitative data in response to it.

And this skepticism becomes more justified when we consider the fact that Nahmias et al., in developing and administering their surveys for this supercomputer study, captured some qualitative evidence suggesting that participants might misinterpret what they are being asked. While in the chapter I’ve been discussing this supercomputer study is presented just as a quantitative study, it turns out that in the more detailed description of this study (which is pointed to in an endnote in the chapter) we learn that participants were invited to explain their responses to their surveys on the backs of the surveys themselves.⁴¹ In the more detailed description we also learn that in pilot studies, participants gave responses indicating a misunderstanding of the survey—specifically, “some seemed to assume that the scenario is impossible *because* Jeremy has free will, rather than making judgments about Jeremy’s freedom on the assumption that the scenario is actual”⁴²—which is important for our skepticism toward the reliability of quantitative surveys in two important ways. First, it shows that Nahmias et al. saw that participants could easily be tripped up on what they are asking, but appear to fail keep this lesson squarely in mind when they later report their general conclusions about what their work shows about the non-intuitiveness of incompatibilism with a significant degree of confidence.⁴³ They do try to remedy this trip up by asking participants to imagine that the supercomputer case is real, but this hardly takes seems a sufficient measure to anticipate the many ways participants might struggle with interpreting the complicated scenario (which I discuss more later in this chapter section).⁴⁴ That is, such qualitative reports, even if they only cropped up in pilot studies, should have been enough to show Nahmias et al. that survey responses can easily indicate a response that a participant doesn’t hold, as evidenced by the instances the researchers reference wherein participants reported

an explanation different than was intended to be captured by the survey. Second, this particular qualitative explanation happens to point to a broader skepticism that I discuss at more length later, which is that people seem to have such attachment to the reality of free will that it is difficult to conclude based on survey responses that participants even digest thought experiments designed to threaten or probe their judgments on free will. The face value lesson from participants assuming the computer scenario is impossible because Jeremy is thought to have free will seems to be that free will is so metaphysically basic for many people that it is more reasonable to them to deny the possibility of supercomputer prediction than it is to entertain the possibility that free will is an illusion. All of this, then, supports that a not-so-fast attitude toward survey-driven claims about free will is appropriate.

It's worth noting that I am not the only one worried about the skeptical picture I've just painted. Shaun Nichols and Joshua Knobe note that free will experiments by Nahmias, Morris, Nadelhoffer, and Turner, including the supercomputer one, rely on "technical vocabulary (e.g., 'laws of nature,' 'current state')" that subjects might not fully understand, which opens up the findings to challenge.⁴⁵ What is clear as day to experimental philosophers might be perfectly opaque to laypeople; and this example isn't even pointing to philosophy-specific terms of art, such as "categorical imperative" or "metaphysical," but instead to jargon that crops up in science departments, humanities departments, and even everyday conversations. Nichols and Knobe bring up another possible worry, as to whether determinism is sufficiently captured in such scenarios, but they also iterate that they aren't interested in dismissing the findings from Nahmias, Morris, Nadelhoffer, and Turner since they think them valuable.

But the MPEP worries emerging from the supercomputer experiment don't end with specific jargon or with descriptions of determinism, though these are key worries. Below are just

a few plausible explanations of the data from the supercomputer experiment that immediately come to mind. Importantly, these explanations cannot be dismissed based on the data from the experiment since I suggest the data *result* from these plausible explanations rather than distinguish between them. And nor can these explanations be dismissed by reference to the general body of data collected by experimental philosophers unless, unbeknownst to me, they include methodological checks that verify that *this particular* vignette—the one in which a future Jeremy is predicted by a supercomputer to rob a bank—produces the kind of judgment that it is supposed to. Here are those plausible explanations, each one named and described with reference to a general “participant”:

- a. *Participant Views Free Will as Axiomatic*: The details of the supercomputer don’t challenge, don’t affect, don’t persuade, and don’t reveal anything of substance because participant holds free will to be axiomatically true. To the participant, the survey question was essentially read as, ‘Do people act freely?’ to which participant responded, ‘Yes.’ To elaborate: participant lives in the real world—that is, *this actual world we are in*—and believes that her actions in the real world are freely willed (in a way that allows for moral responsibility). Jeremy’s hypothetical scenario takes place in the real world, only in the future. Therefore nothing, not even a predictive supercomputer, can negate the freedom in Jeremy’s world (which *is* the real world, in some meaningful sense, that participant occupies while completing the survey). To elaborate, participant’s conviction that her actions are free, and thus Jeremy’s too, is axiomatic and irrefutably grounded such that nothing can undermine it. (You will know that this is not a far-fetched possibility at all if you regularly introduce to others the idea that humans don’t act freely whereby you regularly hear responses—made with

a certainty akin to Descartes' certainty of his existence in the *cogito*—that more or less reduce to 'Nothing you say could convince me I'm not free.'). Therefore Jeremy's world will too maintain free will since it is derived from the actual world, no matter the topsy-turvy features detailed therein. And if participant's conviction in free will, which is immune to challenge, is only being implicitly challenged by Jeremy's derivative world, even if it is in some sense the same world, then this challenge is even *less* forceful—not that that would matter since the conviction in free will is so fundamentally and dogmatically held—than if a challenge arose in participant's real life experience, for example if she found herself in the presence of such a supercomputer. But even this description overstates what is operating because participant's free will conviction is only *perceived* as being challenged from our philosopher's-eye-view on the situation: we know that determinism is at least widely held to be a challenge to free will and so we read this whole research process as a person's conviction being buffeted by an actual challenge, albeit a weak one. But this unique, distant-from-what-participant-experiences-because-participant-isn't-immersed-in-thought-experiments perspective misses what participant is experiencing. The supercomputer details become irrelevant to her decision; she is not appreciating the reality of determinism at all, and is thus not providing a survey response that can be thought to support compatibilism or incompatibilism. She is simply articulating that people, real or otherwise, act freely (in a way that preserves moral responsibility). And so, in sum, participant responds not by expressing a judgment that can be seen as supporting incompatibilism or compatibilism. No, it is far less philosophically interesting (for the goals of the survey) since she doesn't appreciate the role of

determinism; she simply articulates a firm belief that free will is actual for her and for Jeremy. She might actually be a compatibilist, but that wouldn't be revealed unless participant based her survey response on an appreciation of the truth of determinism, and it is unclear how we can know whether participant is responding to an appreciation of the truth of the complicated idea that is determinism. Here, participant judges nothing on behalf of determinism.

- b. *Participant Distinguishes Between Prediction and Determination*: Because the computer is predicting rather than determining, participant thinks that free will still operates—i.e., freedom is contradicted by determination, not by prediction. Knobe and Nichols, also talking about the supercomputer scenario, note something similar: “The story of the supercomputer focuses on the predictability of events in the universe, and many philosophers have taken the predictability of the universe to be less threatening to free will than causal inevitability.”⁴⁶ Therefore, the response participant makes, given this explanation, doesn't capture a judgment regarding compatibilism or incompatibilism. Just as in the previous explanation, participant isn't responding on behalf of determinism as such, though in this case it is because predictability doesn't represent the notion of determinism that is expected to challenge free will convictions. Interestingly, in the more detailed publication of this study (“Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility”), Nahmias et al. report that some participants found the supercomputer example to be impossible for reasons such as “the computer could never acquire that much data, that people could undermine the predictions by learning about them, and that chaos theory or quantum theory makes such predictions impossible.”⁴⁷ To account for this, Nahmias et al. asked participants

to address the scenario on the basis that it *is* actual, but this hardly seems like a sufficient solution given research on the challenges posed by imaginative resistance (which I will discuss in more detail later in this chapter section), which suggests that some people in some cases can't or won't imagine certain things.⁴⁸ This shows the researchers even had evidence in participant qualitative responses to indicate that the quantitative responses might not be so straightforward. While Nahmias et al. grapple with the challenges posed by conveying the right “amount” of determinism, both in the chapter⁴⁹ and the detailed publication⁵⁰, we just don't see enough appreciation of the skepticism toward purely quantitative survey results in relation to such thought experiments when Nahmias et al. conclude that their findings across multiple studies “suggest that most laypersons do *not* have incompatibilist intuitions, though this preliminary work should be supplemented in order to get a firmer grasp on the relevant intuitions.”⁵¹ If anything, I think their work suggests just how difficult it is to get at people's intuitions regarding free will across complicated cases.

- c. *Participant Doubts that Prediction Can Ever Be 100% Accurate*: Participant thinks of prediction as necessarily imperfect and any attempt to predict the future is subject to error. Even after reading that the computer is 100% accurate, participant still thinks of the computer *not* as 100% predictive and therefore Jeremy's choice is not *unfree*. Participant's survey response, then, isn't made on behalf of a philosophical appreciation of determinism, but instead on behalf of an imperfect supercomputer's imperfect attempt at prediction. Thus, participant's response isn't indicative of compatibilism or incompatibilism, as such judgments would require an appreciation of a fully binding determinism, not the imperfect binding of a minutely flawed computer.

There is actually some evidence of this worry about prediction just mentioned above in that participant(s) apparently offered reasons for the impossibility of the supercomputer scenario on behalf of quantum theory's perceived negation of prediction.⁵²

- d. *Participant Thinks Jeremy Still Could Do Otherwise*: Participant understands the scenario perfectly well, but isn't reacting accurately to a notion of determinism, because nothing in the scenario implies strongly enough that Jeremy *couldn't* do otherwise in any given situation. That is, the machine is accurate, it correctly predicts Jeremy's actions, but participant might understand that Jeremy *could* have done otherwise in some sense. Perhaps it would help to create a similar vignette as the supercomputer example that makes determinism's implications more apparent, a vignette that adds the following language: 'When Jeremy robs the bank, he could *not* have done otherwise. Though he had the *feeling* as if he had chosen between alternatives, Jeremy's actions, like all events in the universe, are governed and controlled 100% by the laws of physics: Jeremy robbed the bank and he couldn't have done otherwise. He *did* consciously deliberate on whether to rob the bank, thinking of reasons for and against it, but that experience of deliberation was *completely* scripted by the laws of physics and preceding events and thus had to happen exactly how it did happen.' While this language might read as question-begging to some, I think it more accurately and robustly represents determinism. The original vignette from Nahmias et al. about Jeremy and the supercomputer and an alternative version that includes this language both faithfully describe determinism, so it isn't clear to me why a more anemic description of determinism—like the one given by Nahmias et al.—should be preferable to a comparable vignette that includes my language about Jeremy's actions.

In addition, Nahmias et al.'s anemic description of determinism in the supercomputer example doubtfully alerts people to the fact that they are being asked about a *deterministic* universe, which we have reason to think is likely *not* how they currently view the real world. Drawing from a cross cultural study on deterministic and indeterministic views of our universe, researchers found that, overwhelmingly, people from various cultures hold that the world that we live in is indeterministic, as opposed to deterministic, where the indeterminism applies just to human choice.⁵³ Specifically, subjects in this study addressed vignettes (designed by Nichols and Knobe) that described universes A and B; they were asked which one is most like the universe they live in. Universe A was deterministic through and through, while Universe B was described as follows:

Now imagine a universe (Universe B) in which almost everything that happens is completely caused by whatever happened before it. The one exception is human decision making. For example, one day Mary decided to have French fries at lunch. Since a person's decision in this universe is not completely caused by what happened before it, even if everything in the universe was exactly the same up until Mary made her decision, it did not have to happen that Mary would decide to have French fries. She could have decided to have something different.⁵⁴

The majority of people from different cultures said Universe B is most like the one we all occupy. That is, people largely agreed with the idea that our universe is like a universe in which *everything else but human choice* was determined (a view which saves moral responsibility). This indicates the strong possibility that participant reacted

to Jeremy Hall as a person who acted freely in an otherwise deterministic universe, which doesn't address compatibilism or incompatibilism since such a judgment would be grounded on indeterminism rather than determinism. (Of course, appealing to another quantitative x-phi study to argue that x-phi studies lack methodological justification doesn't get me very far. Nonetheless, I think that participant's thinking in this case is plausible regardless of the study I'm marshaling as evidence—in my experience, people talk about human choices as indeterministically free—and perhaps those who disagree with my claims about the methodological problem *will*—paradoxically—take this referenced study by Hagop Sarkissian et al. to be evidence that the Nahmias et al. supercomputer study lacks methodological justification just because of this plausible explanation of the data which hasn't been ruled out. That is, I don't find this empirical evidence as compelling as others might, but I do find the concern regarding the Nahmias et al. data to be founded nonetheless.)

- e. *Participant Doesn't Understand Determinism Until Its Effects Are Presented:* Participant comprehends the scenario in that no words or phrases are confusing to her, but she, like others, doesn't fully grasp determinism until its *consequences* are very clearly laid out on the table. As a result, her response doesn't indicate support for compatibilism or incompatibilism as both doctrines rely on an understanding of determinism. For example, some participants might need to read about multiple cases in which it *appears* that one can do otherwise, but really her action is determined by prior events. In his article on the x-phi of free will, Tamler Sommers makes a suggestion that speaks to this proposed explanation that participants need deeper education on determinism before providing useful responses on the topic. Specifically,

Sommers recommends that researchers might conduct brief training sessions of 30 to 60 minutes on topics such as free will, determinism, moral responsibility, etc. which allow space for asking for clarification.⁵⁵ This suggestion echoes my worry that determinism is not something fully grasped without some kind of deep education or an expertly negotiated dialectic.

There are potentially more—and more compelling—explanations than the ones I’ve proposed and drawn from, but this list should at least point to the areas of such a survey that are questionable in the MPEP sense: whether or not determinism is appreciated or understood, and whether or not the protagonist’s actions are thought to be appropriately constrained.

Compounding the methodological problem is the fact that no two experimental philosophy projects that employ judgment-capturing surveys are relevantly similar *beyond* their use of judgment-capturing surveys. This is true *even within* the same philosophical debate. Even if experimental philosophers studying free will justify their belief that *their* data from a particular survey based on a particular vignette are reflective of people’s judgments, this does not provide sufficient reason to think that other vignettes relating to free will/moral responsibility produce equally reflective data. Nor does this provide reason that other vignettes relating to other philosophical questions—e.g., those of philosophy of mind, epistemology, etc.—are sufficiently reflective of people’s judgments. What this amounts to then, is the fact that the MPEP is ubiquitous because the quantitative survey is ubiquitous, not because philosophical inquiries are altogether similar. In fact, what falls under philosophy proper is a range of diverse questions that have little in common beyond, say, their complexity, difficulty, and importance to people: questions of meaning, death, knowledge, mind, ethics, aesthetics, existence, language, mathematics, logic, religion, etc.

To better understand the dissimilarity in x-phi projects and x-phi vignettes, let's return to the chapter by Nahmias et al. I will again recount parts of the publication to show how the studies it describes employ thought experiments that are methodologically distinct such that a methodological validation of any one study would not necessarily grant methodological validation to the others despite their apparent similarity. From there, it shouldn't be difficult to see how studies that deal with entirely different philosophical questions—e.g., questions of free will/moral responsibility versus those of epistemology—are methodologically distinct as well. Consider, for example, the dissimilarity between the previous study (supercomputer) and the next one (with a recreated universe); this next one was, by the way, designed in response to worries about the first:

Imagine there is a universe that is re-created over and over again, starting from the exact same initial conditions and with all the same laws of nature. In this universe the same conditions and the same laws of nature produce the exact same outcomes, so that every single time the universe is re-created, everything must happen the exact same way. For instance, in this universe a person named Jill decides to steal a necklace at a particular time, and every time the universe is re-created, Jill decides to steal the necklace at that time.⁵⁶

Here participants were asked if Jill acted freely and whether or not she is blameworthy; 66% of participants agreed with her freedom and 77% to her moral responsibility.⁵⁷ It is easy to see how someone could find these vignettes similar enough such that they think the data from each speaks to the other—after all, both convey ideas of choice and something like determinism. But the devil, or rather the devils in this case, is in the details, as we shall see.

To a degree, Nahmias et al. acknowledge one such devil while missing another when designing this second vignette (the recreated universe). First, they correctly note that, in order to

test the intuitive nature of free will and moral responsibility in response to deterministic thought experiments, “determinism should be as salient to participants as possible without being misleading.”⁵⁸ They fail, however, to detect that what they, as philosophers, think might solve this problem speaks to a deep-seated problem identical or adjacent—I can’t decide which—to the MPEP: philosophers struggle to view thought experiments from the perspective of non-philosophers and thus their attempts to clarify a thought experiment might actually result in further obscuring them. We see this when, following the acknowledgment of the challenge of baking the right amount of determinism into a survey, Nahmias et al. pen a line that is so important to deal with because I think it points right to the MPEP’s root cause, the distance between philosopher and layperson (bold emphasis is mine): “With this in mind, we developed a second scenario **using a simpler...presentation of determinism: Imagine there is a universe that is re-created over and over again**, starting from the exact same initial conditions and with all the same laws of nature.”⁵⁹ This evidences that it is perhaps philosophers’ inability to see the world as non-philosophers do that has produced the MPEP (at least in large part). That is, Nahmias et al. design this second vignette in hopes to simplify the first, they say, and yet they produce something altogether psychedelic.

Let me further explain. It is difficult to overstate how clear a thought experiment can seem to me *once* I’ve entered into it as a room, looked around, picked up its contained objects, and then peered at the world outside from within its confines. And this clarity is in stark contrast with my lack of understanding before I *get it*—I’m unashamed to confess that the thrust of any philosophical zombies⁶⁰ thought experiment was entirely lost on me for some time until it wasn’t. I don’t think I’m unique here, at least not in feeling the qualitative difference between getting and not getting a thought experiment, even if I do take longer to make the transition myself. There seems to be a

large conceptual distance traveled from when we first hear a thought experiment to when we *get it*. I think this helps explain the above quote about simplifying from supercomputers to recreated universes, which, to clarify, is not included so as to be a personal dig but rather as a reminder of the value of x-phi more generally: it's included to remind us philosophers how differently we see the world than non-philosophers. Isn't that why we are asking laypeople things in the first place? (Or at least one reason?) Nahmias et al. produce this statement, without a hint of irony, in reference to their effort to throw something simpler at experimental subjects than a vignette about a predictive supercomputer and a bank robbery that occurs in the year 2195, which also reads as quite weird when you really think about it. (Fidelity Bank is still around in 2195? Banks are still around in 2195? People still have the need to rob banks in 2195? People still have the means by which to rob banks in 2195? Jeremy is still a common-enough name in 2195? And Hall? Predictive supercomputers exist but we haven't squared away bank security in 2195?) But the second scenario, that of the re-created universe, is decidedly more bizarre. And the notion that a scenario in which the universe results from some samsaric act of Brahma-creation satisfies an effort for simplicity evidences that philosophers often occupy a unique headspace.⁶¹ Don't get me wrong; I get it. For me, this *is* simpler and certainly easier to follow than the supercomputer case—we don't have to get into the whole bit about whether prediction is determination. But as a philosopher, I'm also used to introducing multiple universes into everyday thinking and common conversation; laypeople aren't.

To connect this re-created universe study, too, with the MPEP, I want to point to the very different way that determinism is expressed and how it is still quite possible that participants are reacting to things other than an appreciation of determinism's truth: the first scenario relies on a predictive supercomputer and this one involves a re-created universe, yet both run similar risks of

failing to capture their intended judgments. So, when authors note that the results of this second scenario are similar to the previous scenario,⁶² they do so without addressing the still-lurking methodological problem in two important ways. First, this assumes that the results from this second study are grounded on the desired judgments being captured. Second, this assumes that there is some meaningful relationship between the two studies such that the results of one connect with the results of the other; but even if a single survey could be shown to overcome the MPEP, this would not demonstrate that any other surveys, too, would overcome the MPEP even if both surveys are tapping similar philosophical questions. Here then are some potential arguments (similar in form to the plausible explanations above) that might explain the two scenarios in conjunction *other* than the explanation that participants really are demonstrating their judgments regarding free will, compatibilism, and moral responsibility as affected by an accurate appreciation of determinism:

- a. *Participants View Free Will as Axiomatic in Both Scenarios:* Neither the supercomputer nor the re-created universe scenario taps a judgment in response to determinism because it isn't really being considered by participants. Participants hold that free will is axiomatically true, and so their responses in both cases grant free will to the agents in the surveys since these agents occupy universes analogous to our own. Therefore, participants think the acts in the surveys are freely willed and thus the participants aren't reporting judgments on behalf of the truth of or even consideration of determinism.
- b. *Only the First Scenario Succeeds:* The first scenario, with the supercomputer, really does tap into the desired judgment, but the second one, with the re-created universe, does not. In the first scenario, the participant learns that Jeremy "robs" a bank and no

verb implying choice is introduced; participants do hold some kind of natural compatibilism.⁶³ However, in the second scenario, participant learns that Jill “decides to steal a necklace,” which implies that it was an indeterministically made, free choice.⁶⁴ Therefore, the apparent similarity in results is misleading as the explanation for a certain response on the first scenario is completely different than the relevantly similar response on the second scenario.

- c. *Only the Second Scenario Succeeds*: The first scenario, with the supercomputer, doesn’t tap the desired intuition, but the second one, with the re-created universe, does. The predictive supercomputer scenario implies that the future is predicted but not determined, but the second one clearly says that “everything must happen the exact same way.”⁶⁵ And so, again, any apparently similar responses are explained by different causes.

Since neither of these thought experiments has, to my knowledge, been rigorously tested for validity based on multiple methods⁶⁶, it is difficult to see how we can be confident that either really is tapping and capturing the desired judgments. This is obviously problematic when looking at either scenario in isolation as evidentiary support for a claim about laypeople’s intuitions, and even more problematic when the two are combined as evidence of something: one survey might be successfully achieving its ends while the other isn’t, and yet the latter is used in conjunction with the former to *further validate* the findings of the former (or vice versa). There are, then, many problems with using such survey responses to say anything about people’s specific judgments/intuitions about free will and moral responsibility. In sum, even if both scenarios produce similar responses in terms of what *appears* to be compatibilist judgments/intuitions, it is

entirely plausible that one, the other, or both fail to achieve this finding on the basis of genuine compatibilist judgments/intuitions.

Now, with a shorter analysis, we will look at the last experiment in this collection from the Nahmias et al. chapter, which again plays around with the way determinism is expressed. And again, the survey lacks any empirically verifiable methodological justification—the MPEP is back. In this last scenario, the authors make the characters’ “genes and upbringing” the salient forces behind their actions:

Imagine there is a world where the beliefs and values of every person are caused completely by the combination of one’s genes and one’s environment. For instance, one day in this world, two identical twins, named Fred and Barney, are born to a mother who puts them up for adoption. Fred is adopted by the Jerksens and Barney is adopted by the Kindersens. In Fred’s case, his genes and his upbringing by the selfish Jerkson family have caused him to value money above all else and to believe it is OK to acquire money however you can. In Barney’s case, his (identical) genes and his upbringing by the kindly Kinderson family have caused him to value honesty above all else and to believe one should always respect others’ property. Both Fred and Barney are intelligent individuals who are capable of deliberating about what they do.

One day Fred and Barney each happen to find a wallet containing \$1,000 and the identification of the owner (neither man knows the owner). Each man is sure there is nobody else around. After deliberation, Fred Jerkson, because of his beliefs and values, keeps the money. After deliberation, Barney Kinderson, because of his beliefs and values, returns the wallet to its owner.

Given that, in this world, one's genes and environment completely cause one's beliefs and values, it is true that if Fred had been adopted by the Kindersons, he would have had the beliefs and values that would have caused him to return the wallet; and if Barney had been adopted by the Jerksens, he would have had the beliefs and values that would have caused him to keep the wallet.⁶⁷

In response, participants were asked about the decisions of Fred and Barney—76% of participants “judged both that Fred kept the wallet of his own free will and that Barney returned it of his own free will.”⁶⁸ Again, we note that this last study also displays a completely different expression of determinism, this one focusing more on a classic psychological account of behavior: that it's the product of genes and environment.

While the researchers have designed this so as to capture people's intuitional responses to determinism as it relates to free will, I still find a number of other plausible explanations for the data of this study in isolation and in conjunction with the others—the MPEP is still here. I will, though, just present the one that I think holds the most weight: many people have such a deeply ingrained belief that humans have free will—i.e., the ability to do otherwise in a given situation that preserves moral responsibility—that no matter how one tries to express determinism in a brief thought experiment—by pointing to predictive computers, re-creating universes, the effects of genes and environment, or other clever constructions—participants will more often than not *still* ascribe free will to actors in a given “world” because any “world” they are asked to imagine will be based on the world they know in which they live, the one where actions are thought to be freely willed. In brief, participants view free will as axiomatically true in any kind of “world” that is essentially like ours.

And now I want to drill down on this one objection, this one plausible explanation that might explain much of what Nahmias et al. are finding. In actuality, I *need* to drill down on this to justify my MPEP worries. Sure, there are many alternative explanations for a given data set that might plausibly explain the data better than the explanation proposed by the researchers themselves. But the researchers might find each alternative explanation only moderately worrying—‘Yes, this or that alternative explanation *could* be what’s going on,’ they might think, ‘but probably not; this alternative explanation isn’t that threatening in the end.’ Or, an even better defense might note that while there is a potential flaw in each survey designed to test subjects’ judgments, each survey is flawed in a *different way* such that the story told by the many studies considered in total is that subjects really do think the way predicted by each flawed study: ‘Yes,’ the researchers might conclude, ‘study 1 might underexplain this philosophical concept; study 2 might overexplain it; study 3 plausibly can be interpreted three different ways; and yet they each produced results supportive of theory *x*. The commonality here is support for theory *x* which is best explained by *x* being true.’ And this might just be the case for these studies on compatibilism by Nahmias et al. The supercomputer is plausibly flawed for different reasons than the re-created universe which is plausibly flawed for different reasons than the case of Barney and Fred; and yet compatibilism is the one theory that unifies them all. Without an alternate explanation that unifies the data from these multiple studies, Nahmias et al. might be right to view the convergence in these studies—flawed as they might each be in idiosyncratic ways—to be supportive of their conclusion: through abductive reasoning, prevalence of compatibilist lay judgments really is the best explanation.

I argue, though, that we do have, in addition to good alternative explanations for any *one* study by Nahmias et al., a good alternative explanation that unifies these studies in a way so

fundamental that their individual flaws can be viewed as secondary. Before elaborating on the specifics, my proposed explanation is, in general, that some people view agents as entities that have free will. This explanation is based on a kind of judgment similar to what Feltz and Millan call *free will no matter what* intuitions, which are (at least⁶⁹) “intuitions indicating that fated actions can be freely done and one can be morally responsible for them.”⁷⁰ Based on their research, they propose that at least some people who may appear to be natural compatibilists are really just hardline free will ascribers, even in cases of an agent’s being fated to do something.⁷¹ I say similar, because the judgment I’m concerned with isn’t necessarily tied to fatalism.⁷² What I’m thinking of can be described as a general tendency to ascribe free will (a robust kind that preserves moral responsibility) in most cases where an agent is acting (but not *no matter what*). (I’ve suggested up until this point in this section that subjects in the studies discussed so far might hold free will to be axiomatically the case when they are asked about a world relevantly similar to our own. And here again I am talking about this kind of axiomatic view of free will, though I will be taking the idea and clarifying it a bit more; the axiomatic notion discussed previously, though, is the basis for this current analysis.) Rather than free will no matter what, I’ll call this *if possible then actual free will* (IPTA), because it seems to be the case that some subjects preserve an agent’s free will so long as they view it as possible that agent acts freely, regardless of whether it really is possible. A subject’s assessment of the possibility of free will—which I think leads directly to a conclusion of the actuality of free will—could result from reading a vignette that treats free will in a given case as an open question, or it might result from their own assessment of a scenario such that it seems like free will *might* be possible, even if it isn’t. What appears to explain this kind of judgment is the strong, deeply held view—an axiomatic one—that many people have that persons are agents that act freely, full stop. Another way we might think of IPTA is that it is a very stubborn, mile-

from-an-inch kind of free will: whenever it seems there is a little bit of wiggle room in a scenario that leaves open the possibility for free will, or some possible understanding of a situation that could preserve free will, a subject will swing open the door left slightly ajar and ascribe free will to someone. These subjects certainly don't need to be talked into someone having free will. For them, free will (and in conjunction, moral responsibility) is metaphysically basic.

When some people hear about a possible challenge to free will, it might sound like the philosopher is in all earnest approaching them to say, 'Let me tell you why you're actually not alive right now.' In hearing this, the layperson, I imagine, must conclude that the philosopher is playing with words in some obnoxious way, redefining "alive" or "dead" or something else. She might stick around to hear what the philosopher has to say, but there is no real breath-holding for the punch-line; the secret at the end of the lecture is not going to fundamentally challenge her belief—her *knowledge*—that she is alive. And I think, for at least a significant portion of laypeople, this might be how they view free will: 'Give me your clever argument, philosopher, but I'm still *free* to roll my eyes at its conclusion.' And this free will, the IPTA kind, seems to explain these findings by Nahmias et al. across multiple studies just as well, if not better, than the notion that laypeople have some kind of sophisticated theory like compatibilism—holding a firm understanding of determinism and free will and judging the two to be metaphysically reconcilable. Importantly, I defend the plausibility of this explanation, based on IPTA judgments, not because I aim to resolve to the free will debate, but because I want this example to evidence the MPEP. If an alternative explanation like IPTA judgments unifies a series of studies on a single topic *and* is different than the explanation proposed by the original researchers, then it stands to reason that we lack justification in accepting one or the other explanation. And, because this stalemate appears to be borne out of an over-reliance on quantitative surveys—an over-reliance evidenced by the fact

that the surveys themselves aren't able to resolve the current debate—then we are back at the MPEP in need of new methodological approaches to shed light on the problem from new angles.

Of course, I've been pointing at something like the IPTA all along as a plausible alternative explanation for this or that experiment by Nahmias et al. But here I'd like to argue that it plausibly cuts across each study in a deep way and therefore isn't merely a problem for this or that particular experiment. This unifying explanation is worrisome because it is grounded on what appears to be a powerful intuition: that humans have free will. And, because it plausibly cuts through multiple studies, it can't be dismissed by an appeal to one or multiple studies that don't account for something like it (the IPTA judgment)—in a sense, it is a more problematic version of the MPEP. In its more basic form, the methodological problem rears its head because experimental philosophers aren't regularly providing justification that their individual surveys capture their desired judgments. But the problem grows increasingly severe when we have good reasons for thinking that such method checking would reveal that the desired judgments *aren't* being captured by any one study *or* some more total picture revealed by their combination. And, in the free will debate, the good reason we have for suspecting that we haven't captured compatibilist judgments is the apparently stubborn judgment that humans *just have* free will, a judgment that many people plausibly hold.

But this might look like speculation; I've simply asserted it plausible that people have this strong free will judgment I call IPTA. Here I'll show this is plausible beyond mere speculation, and I'll show why this is a powerful version of the methodological problem. And just to nip a future objection in the bud, it might seem like what I am achieving in arguing for the plausibility of this uniquely powerful methodological problem is showing that the case of free will is unique in philosophy. In some sense this is true, the problem of free will is a problem *because* of one of

the strongest judgments/intuitions humans appear to have—that they act freely. However, I argue that the common debates experimental philosophers find themselves wading into have similarly basic, primary, and dogmatic judgments as part of the story: that is in some sense what defines a branch or ongoing debate in philosophy. We talk of issues in ethics, metaphysics, the philosophy of religion, epistemology, the mind-body problem, and so much more because we run up against certain basic judgments/intuitions that don't square easily with other judgments/intuitions, empirical facts, cultural traditions, etc. So, the problem of strong, basic judgments/intuitions isn't unique to free will in any real sense. In other words, we return time and time again to our barns and trolleys—and the more expansive philosophical traditions that they respectively encapsulate—*because* there are not universally agreed upon basic judgments/intuitions that can settle the debates they embody via cleverly constructed thought experiments. Without further ado, let's look more closely at an especially powerful version of the methodological problem as it plausibly manifests in the x-phi of free will as a result of what seems to be a common and strong judgment/intuition: that we humans act freely. As long as this judgment is a plausible explanation of certain x-phi of free will data, it warrants our investigation.

To begin, I want to note that I was happy to concoct something like the IPTA theory prior to discovering the similar worry brought about by Feltz and Millan (about no matter what intuitions) since our like thinking validates, to some small degree, my (and their) sense that some people just think actors acting have free will no matter what (or close to it). Again, from their research, Feltz and Millan claim that “many people judge that a person has free will no matter what” such that they “judge that a person freely performs and is responsible for a fated action.”⁷³ This is noteworthy, because while we might expect people to ascribe free will in determined cases or predicted cases, fated cases are just the kind we imagine people will say removes the possibility

of free will. If further research corroborates their findings, then we have good reason to think the intuition that actions are things acted upon based on free will runs deep. I don't want to lean too heavily on Feltz and Millan's claims since I am systematically casting doubt on data that is derived largely from quantitative surveys, as is the case with their data. (In fact, below, I'll bring up some skepticism toward the kind of strange vignettes they employ as part of their project.) But *that* my personal experience talking about and thinking about free will led me to something like the IPTA theory before I saw Feltz and Millan's work, which too supports some like the IPTA, counts as *some* kind of evidence that deep-rooted free will judgments/intuitions plausibly explain what Nahmias et al. took to be compatibilist intuitions.

Of course, my promotion of the IPTA theory too might be inaccurate. Andow and Cova, for example, push back on Feltz and Millan's conclusions about free will no matter intuitions.⁷⁴ But this rebuttal also points at the MPEP, again arising from an over-reliance on quantitative surveys: Andow and Cova base their disagreement with Feltz and Millan on surveys of their own. And if we doubt that any set of surveys alone—all of which depend on certain methodological limitations—can reveal the truth of laypeople's judgments, then it is unclear how *more* surveys can tell us which proposed explanation of laypeople's judgments, the one by Feltz and Millan or by Andow and Cova, is preferable. Still, Andow and Cova argue on the basis of their findings that one explanation is better than another:

So, we hypothesized that Feltz and Millan's results are better explained by supposing not that many of their participants had [free will no matter what] intuitions, but rather that many had broadly compatibilist intuitions—source compatibilist intuitions—and had read their fatalistic cases as allowing agents' mental states to be the source of their actions. We

tested our hypothesis through two studies, the results of which supported our interpretation of Feltz and Millan's results.⁷⁵

But, when Andow and Cova disagree with much of Feltz and Millan's thinking and hold that some participants might not be ascribing free will no matter what, but rather are ascribing free will as long as people are the source of their actions (i.e., some participants are "source compatibilists"⁷⁶), I don't think we are necessarily closer to finding what explanation of Nahmias et al.'s data is right. At this point, I feel justified in being a bit dismissive because I think the data points to a methodologically defined stalemate since Andow and Cova base their counter to Feltz and Millan on thought experiments that are even more bizarre and even less persuasively revelatory of people's potentially complex judgments. Consider the following vignette they employ, More Extreme Book, which is derivative from vignettes used by Feltz and Millan:

Imagine that there is a magic book that has all of our decisions and actions truly written in its content. For instance, whenever we are trying to decide what to do, the decision we end up making is completely and truly written in this book and the decision will happen regardless of our thoughts, beliefs, desires, and plans. The magic book has these events truly written in it lifetimes before the events took place. So, if the book has an event written in it, the event will definitely occur regardless of the past events and the laws of nature. Even if a person does not want to act this way, then she will be forced to act against her will by the book's magical powers. For example, one day a person named John decides to kill his wife so that he can marry his neighbor and he does it. Deep down, John did not want to kill his wife, and did not want to marry his neighbor. However, since these specific events were already and truly written in the book, it was impossible for John not to kill his

wife regardless of his thoughts, beliefs, desires, and plans. Assume the book's contents made it impossible for John not to kill his wife.⁷⁷

Here, the debate over which nuanced or non-nuanced judgments/intuitions participants have regarding free will has gone off the rails with no compelling empirical evidence showing that one or another elaborate vignette—those involving fatalistic magic books⁷⁸ to source-compatibilist-incompatible, fatalistic magic books⁷⁹—reveals what most participants judge/intuit regarding free will, determinism, fatalism, etc. I think Feltz and Millan put it best when they say what I take to be a kind of vote of confidence for qualitative methodology to help out in x-phi:

Almost all of the scenarios used in the experimental philosophy of free will have face validity.... But face validity may not be sufficient to settle many of the issues involved in the experimental philosophy of free will... Pure survey methodology in some experimental philosophy has perhaps run its course. Rather, using other methodologies that get closer to the proximal judgment processes involved would be more effective and compelling.... For example, protocol analysis that pays particular attention to individual differences is likely to yield better, higher-fidelity data, that could help weigh in on some of these issues.⁸⁰

And that's really the point: who can truly say, after all of the quantitative surveys, *what* people even judge/intuit regarding determinism and free will given that subtle word changes bring about different results and there are no mixed method approaches to ensure that participants are even reacting to the *right things* when they provide their responses. And we don't even have to get to these magic books before we see a problem. We find ourselves in the morass of the complexity of philosophical debates on free will when Feltz and Millan note one challenge of baking fatalism into vignettes for experimentation is that "there are a number of possible ways to describe fatalism."⁸¹ That's right; there are. And if these concepts are this controversial to philosophers,

who spend their time coming to understand doctrines such as fatalism, it is hard to see how we can view ourselves as cutting through the debate and getting to the bottom of knowing laypeople's judgments/intuitions. For just these reasons, my study (which I will detail later) actually incorporates protocol analysis of a kind in accordance with what Feltz and Millan suggest.

Ultimately, I'm eager to get to my study for some of the reasons I've been highlighting through this free will/moral responsibility research, but I realize there is a great deal that needs to be set in place before I do so. Even this chapter has been a lot of things centered around the MPEP, and while we're close to wrapping up, a few more things need to be said. Specifically, Nahmias et al. reply to some potential objections to their work, which I want to highlight since I've chosen their chapter as a focal point for *my* point about methodology. I want to note an important theme in the objections they deal with in this chapter and, more importantly, point to two objections they don't deal with. I argue that for both the important theme in the objections they deal with and the objections they don't, the overarching lesson to be learned is that an over-reliance on quantitative surveys is problematic and limiting.

The important theme in the objections they *do deal with* can be summed up by the worry that determinism wasn't adequately or fully or accurately described in the vignettes. However, I don't think Nahmias et al. provide a satisfying response to this general objection, and instead they show just how difficult a position the x-philosopher is in who wants to use surveys to measure laypeople's judgments/intuitions on free will. For example, Nahmias et al. give this response to a potential objection: "One might argue that most people will *not* in fact recognize a conflict between determinism and freedom or responsibility, but will only come to see such a conflict once they understand the implications of determinism. Of course, 'getting people to see these implications' is probably going to be a euphemism for 'giving them a philosophical argument,' and an

incompatibilist one at that.”⁸² The problem is that this points to a dilemma for the experimental philosopher as a result of her wanting to capture a layperson’s natural, unperturbed, untaught, unmolded judgment/intuition regarding determinism and its implications for free will without her vignette contaminating the judgment/intuition by way of serving as a lesson *on determinism* which might affect the pre-theoretical judgment/intuition that existed prior to such teaching. One horn of the dilemma is administering a vignette that adequately details the constraints of determinism, which appears to amount to some level of ‘teaching’ determinism, which likely adulterates the judgment/intuition. The other horn is that abstinence from experimentation altogether, which grants no access to laypeople’s judgments/intuitions on free will. This dilemma results because determinism itself is a complicated doctrine such that the kind of tool needed to measure its influence on a person’s judgments/intuitions on free will appears to be a decently invasive device: a survey question in which some of the core details of and implications of determinism are explained and, in a sense, *argued for*, as Nahmias et al. fear.

The way I see this dilemma, then, is that the x-philosopher is in a very difficult position in wanting to get at people’s judgments and will find no solution in relying solely on quantitative surveys. While it may be impossible to learn about a person’s thinking on some philosophical matter without contaminating their judgments—i.e., it might be impossible to extract a judgment without altering it to some small degree—the solution to this problem isn’t to keep administering minimally invasive vignettes in the hopes that they are accurately capturing complex judgments when they might only be picking up noise. One solution, then, might be to work with some people in a more methodologically in-depth manner, such as a qualitative interview process, in order to get a glimpse at precisely how they are understanding the concepts in the vignette so as to revise the vignette for later employment with other people. Nahmias et al. themselves note that

interviews in particular might be a good solution to understanding people's judgments.⁸³ Qualitative methods might not solve the dilemma altogether by negating the effects of the first horn since they too involve affecting a person's thinking, but they can potentially shed light on ways to more accurately capture judgments while influencing them less in the process. What I'm imagining here is a multi-study process whereby the first study involves qualitative methods that will hopefully reveal insights into how people interpret a given vignette so that the next study can employ quantitative surveys that have been revised based on the insights drawn from the interviews in the first study. For example, we might get two very different reactions in interviews when discussing the supercomputer vignette and the re-created universe vignette. Let's imagine that qualitative interviews reveal that subjects can consistently talk about the re-created universe vignette in a way that sufficiently understands determinism and the interviewer finds herself rarely needing to intervene to correct subjects' notion of determinism on its behalf; this tells us that this vignette gets people to get determinism quite effortlessly and with little invasiveness. In addition, let's imagine that interviews about the supercomputer example reveal that people aren't getting determinism from it, because they disbelieve the computer's predictive power, and therefore need further teaching and intervention until they get determinism. From this hypothetical data, it would be clear that the re-created universe gives us better access to subjects' black-box judgments regarding determinism because it elicits an understanding of determinism without further teaching; researchers could then opt for making the re-created universe example the canonical approach to getting people to react to determinism, over the supercomputer. Of course, subjects in these interviews have already been exposed to too much concept-contamination or teaching such that their views on a subsequent survey might be too affected to be viewed as reliable, but the insights these subjects produced could influence surveys handed out to countless *other* people for the sake

of x-phi research. In the end, I think Nahmias et al. are right to worry about subjects getting argued a certain position, specifically a position their objectors support, but I think their worry points to something deeper than a debate between incompatibilism and compatibilism. It points to the problem with placing so much stake in quantitative surveys alone to clarify and resolve debates between experimental philosophers.

Now to the objection they don't deal with: in the objections that Nahmias et al. deal with, none discuss the fact that none of their vignettes and questions, to my knowledge, achieve anything beyond face validity. The researchers design what they think will tap a certain intuition, they make some changes and run it again, make some changes and run it again. Each time, though, there is nothing that the researchers can point to for evidence that any one of these vignettes—i.e., with the supercomputer, with the re-created universe, with the twins separated at birth—is valid. At most, these vignettes have face validity—i.e., they appear to capture that they are intended to capture—but even that is suspect. If we return to the start of the previously mentioned quote from Feltz and Millan—“Almost all of the scenarios used in the experimental philosophy of free will have face validity”⁸⁴—we might be tempted to nod in agreement because this generally seems right. However, I think we should be suspicious of even granting face validity to so many vignettes. While I agree with the sentiment expressed by the larger point Feltz and Millan make in the full quote that this is excerpted from—i.e., that it's difficult to rely on surveys alone to get to the bottom of laypeople's judgments/intuitions—I think that this initial claim about face validity might be an overstatement of what x-phi surveys achieve. If by face validity we mean *viewed as valid by philosophers who view clever thought experiments as effective tools to pump genuine judgments/intuitions*, then yes, many x-phi thought experiments have face validity. But if we worry that these thought experiments aren't so effective, then this face validity claim isn't

compelling. It really might be that a claim of face validity says more about the researcher ascribing such validity than the thought experiment itself: if someone views the magic book construction as a viable way to get at someone's judgment/intuition toward free will, it might just be that *she* views the magic book as a clear articulation of certain manifestations of her perception of the free will debate. In addition, it is worth noting that philosophers might be in the worst position to judge what is or isn't valid on its face given that there appears to be a significant distance traveled when one trains as a philosopher such that they can read outlandish thought experiments as perfectly sensible scenarios when laypeople might be fundamentally perplexed by them. Again, I see the lesson to be learned here is that an over-reliance on quantitative surveys is problematic and limiting.

Now, one might argue that the current debates surrounding Nahmias et al. *serve* the function of *testing* for an appropriate measure of intuitive compatibilism or incompatibilism that goes beyond face validity. That is, the process of continued research and discussion functions as a kind of development of effective surveys: first philosophers review Nahmias et al.'s findings, then they develop their own thought experiments and surveys in response, and next come up with new data, then others respond to these data in conjunction with those from Nahmias et al., etc. However, Nahmias et al. do not present their work as the development of a valid measurement of free will intuitions *and* they make claims based off of their results rather than remain agnostic until better methods develop, all of which indicates that their data is being treated as substantive, not merely as a kind of un-validated test of free will that needs undergo significant replication and multi-method validation before it can be thought to produce substantive data. Here is such a claim that indicates the data are taken to be substantive: "The results from these three studies...offer considerable evidence for the falsity of the incompatibilist prediction (P)—i.e., the predication that

most ordinary people would judge that agents in a deterministic scenario do not act of their own free will and are not morally responsible.”⁸⁵ Later, the authors do hedge quite a bit and treat their results as preliminary, but this is quite different than saying that the results ought *not* be taken to be substantive *until* further validation occurs.⁸⁶ Their total project rings of the view that these data are substantive unless and until further studies show that they are not: they place the onus on future incompatibilists and future researchers as if they currently have provided a successful volley to other thinkers.⁸⁷

And there is one last objection Nahmias et al. don't deal with in their chapter which also points to a need for more than quantitative surveys: it is unclear if compatibilism as embodied in their surveys is a coherent philosophical position in and of itself that laypeople can speak to such that it becomes problematic to use people's endorsements of a potentially incoherent view (i.e., compatibilism) as evidence that they have *non-incompatibilist* judgments. That is, I worry that the evidence *against* incompatibilism in these studies is participant endorsement of a position (i.e., compatibilism) that is potentially incoherent. For an incompatibilist, free will really is incompatible with determinism such that it might be incoherent or contradictory to claim or endorse that compatibilism is the case. And if this is the case, then it is difficult to know if people selecting compatibilist-friendly (or, strictly speaking, *non-incompatibilist*) responses on quantitative surveys, such as Jeremy acting freely in robbing the bank, are fully understanding what they might be viewed as subscribing to. Or, to put this another way, I'm not convinced we can confidently label someone a compatibilist (or claim he or she has compatibilist judgments) or a *non-incompatibilist* (or claim he or she has *non-incompatibilist* judgments) based on her endorsement of compatibilist-friendly (or *non-incompatibilist*) survey responses given the potential incoherence of compatibilism. I do think I'd be convinced if I saw that she can articulate

some complex, unique claims *in her own words* about compatibilism and incompatibilism, because this would show me that she has a sufficient grasp on what she is endorsing such that she is able to articulate the complex view herself. There might be other methods, too, that we can use to clarify if someone is a compatibilist (or a *non-incompatibilist*), but I imagine they would need to be more sophisticated or open-ended than the quantitative survey questions typical of x-phi research. To be clear, Nahmias et al. are aware that a participant's rejection of incompatibilism doesn't necessarily mean that she holds a compatibilist view—this is evident in the conclusions quoted just above from their chapter about laypeople not having incompatibilist judgments as well as in the detailed publication wherein they specifically state, after discussing the supercomputer case and the identical twins case, that “by claiming that most people do not express incompatibilist intuitions in these cases, we are *not* endorsing the stronger claim that the folk *do* have compatibilist intuitions.”⁸⁸ Nonetheless, it is completely reasonable to view participants' compatibilist-friendly (or *non-incompatibilist*) responses to the supercomputer case, the re-created universe case, and the identical twins case as compatibilist judgments if one holds compatibilism to be a coherent notion: participants are presented with some kind of description of determinism and they end up endorsing free action in spite of it as if such a notion is coherent. But the problem with this is that compatibilism, as implied by responses that endorse free action in spite of determinism, might *not* be a coherent position and therefore it is difficult to know how to make sense of participants who opted for the compatibilist-friendly response to the cases. If compatibilism really is incoherent, then it's just difficult to understand how someone could genuinely hold it to be true in any scenario let alone after only reading some frighteningly brief introduction to determinism in a survey. That said, I would be willing to put a hold on my deeper worry about compatibilism (or *non-incompatibilism*) being an incoherent position if there were good evidence that participants came

to embody and articulate decently coherent-sounding judgments during, say, a qualitative interview on multiple vignettes pertaining to free will. For example, if subjects regularly uttered sentences like the following, I'd be more inclined to concede that they had compatibilist judgments, or at the very least *non-incompatibilist* judgments: 'Yes, Jeremy will rob the bank, and there is no other course of action that could happen since all of his actions and thoughts and wants are determined by prior events. However, I really think that Jeremy still has a free choice because he is acting on his desire to rob the bank and if he had desired to not rob the bank, he would not have robbed the bank. Of course, his want *was* to rob it, and it couldn't have been otherwise; he was determined by prior events to have that want. His life experiences and genes lead him to want to rob the bank but he still wanted to and no one was forcing him to or threatening him to, and that's freedom in my view.' Without this kind of understanding—or at least something that approximates it—it is difficult to know whether subjects are reporting compatibilist (or *non-incompatibilist*) judgments or something that merely appears as such.

In fact, if compatibilism *is* incoherent, then survey results showing subjects hold compatibilist judgments can be interpreted as invalid in some sense. How could people be reporting that they hold an incoherent view? To sharpen up my reasoning here, consider how we might interpret results from a survey that presented subjects with a claim that is more commonly accepted to be incoherent, namely the doctrine of the Trinity, which “is commonly expressed as the statement that the one God exists as or in three equally divine ‘persons’, the Father, the Son, and the Holy Spirit.”⁸⁹ Not concerning ourselves with the many interesting ways this doctrine has been and is interpreted by Christian theologians, we can create a thought experiment that relies on a version of it which most clearly reads as incoherent: *Imagine there is a divine supercomputer that only reports true statements. That is, every output from this supercomputer is 100% true.*

*One day, the supercomputer reports that the doctrine of the Trinity is true, God is and always has been exactly one person and exactly three persons. Is the doctrine of the Trinity true?*⁹⁰ If subjects overwhelmingly report the doctrine is true, then it would be difficult to know how to interpret the results. Are they simply *going along* with the scenario and just affirming the truth of anything the supercomputer is said to report based on the description that it is always correct? If so, what does this mean for their judgments about the truth of the Trinity outside of this scenario? Or, if subjects overwhelmingly respond that the doctrine is *not* true based on the scenario, does this mean that they reject the infallibility of the computer? Or are they just rejecting that three and one could ever be identical? Or might responses reflect something different altogether? The data might reflect issues of imaginative resistance, which, in general terms, “occurs when an otherwise competent imaginer finds it difficult to engage in some sort of prompted imaginative activity,”⁹¹ since it could be that subjects outright resist the premise of the supercomputer that grants truth to the doctrine of the Trinity. In the end, subjects’ responses might be somewhat strange in that they don’t track this or that theory of the Trinity but rather reflect just *some* kind of response in the face of an incoherent survey prompt. I think it safe to say that results from such a hypothetical survey wouldn’t be taken to indicate much of anything, at least not on their own. And if we wouldn’t find ourselves comfortably making claims about people’s judgments based on this Trinity survey, then I think the same reservation is appropriate in the face of surveys that present compatibilism as a plausible doctrine since compatibilism might be incoherent.

The solution to my worries about Nahmias et al.’s claims, including the one of incoherence of compatibilism, might be to do away with as much of the restrictions of the quantitative survey without destroying a given thought experiment of interest; I see qualitative methodologies helping here. In a qualitative response, such as during an interview, but not in the typical quantitative

survey, a subject can challenge the confines of a trolley problem scenario or seek clarification as to just how binding the predictions of a supercomputer really are. If they seem to think compatibilism is incoherent or incorrect, we can get a sense of that reality in their utterances if they reflect confusion as to how a person could freely act in light of determinism. Allowing for subjects to clarify creates a feedback loop, then, between the subject and researcher. For example, some subjects might not know what is meant when the vignette says the computer “looks at the state of the universe at a certain time”⁹² or what it means that the computer “deduces”⁹³ certain things; in a follow up interview, this can be clarified and later the survey revised if necessary so as to avoid future confusion. Currently, it seems the procedure is that experimental philosophers design a vignette based on what *they* think it will prime in the minds of their subjects, they then collect and review data, read the responses of other philosophers and psychologists, and then make revisions that *they* think will clarify any instances of ambiguity or unclarity in the vignette. But this process leaves out any rich feedback from the subjects themselves who are truly the judges of how the surveys are interpreted. Qualitative methodologies can remedy this, because they allow for subjects to provide feedback in addition to their quantitative survey responses.

In addition, the layer of qualitative data from an interview can point to the formation of relevant or irrelevant judgments. If one by one subjects note that they find a predictive supercomputer to be unrealistic, and yet responded to the survey anyways, then this seems to convey that determinism as described by the supercomputer isn't being fully appreciated by some subjects. Or, if subject after subject describes the scenario in fatalistic rather than deterministic terms based on a misunderstanding of the intention of the supercomputer vignette—e.g., ‘Why would anyone program a computer so someone has to rob a bank? That seems wrong’—then it is reasonable to worry that some subjects aren't tracking the key plot points of the survey and revision

might be necessary. Also, with qualitative methodologies, subjects can highlight reactions they had to the survey that they might have taken to be unimportant but are in fact philosophically interesting. For example, we could imagine that multiple subjects, during a follow-up interview after the survey, report that they were initially struck with the thought that today's computers aren't that far off from predicting people's behaviors like the supercomputer did in the vignette. Clearly, this reaction is interesting for the topic at hand, and might help us to understand what kind of judgment subjects were making. It still might not be clear if this person's judgment turns out to be compatibilist or incompatibilist, but this information might lead us to think that the supercomputer in the vignette looks too similar in its computing power to computers available today in the real world such that subjects might not realize how truly alien the hypothetical supercomputer is and how alien its potential effects might be on how we view decision-making. Based on this, researchers could craft a different vignette which didn't run into this potential snag. Certainly, there are many more benefits that qualitative methodologies can bring, and I will explore those as this entire project unfolds. In their detailed publication ("Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility"), Nahmias et al. too appreciate some role for qualitative feedback in that they are sensitive to things participants write on the back of their surveys and later recommend the value of interviews.⁹⁴ For now, it should be clear that opportunities for subjects to seek clarification, to push beyond the boundaries of the thought experiment, and to report additional reactions to a survey all provide a healthy feedback loop of information not only on a given survey, but on the philosophical topic being explored in a given research project.

In the end, I don't think the challenge is one of discovering people's judgments/intuitions regarding issues connected to ethics and beyond, but one of discovering how to discover people's

judgments/intuitions. Were we confident that we have the precise means by which to extract sophisticated judgments/intuitions from laypeople, then I would think differently, but determinism and compatibilism and free will find themselves as the center of heated and extended debates because they *aren't* straightforward ideas. To sum up with regard to these studies, then, I'll note my alignment with Sommers, who references experimental philosophy work on free will including that from Nahmias et al., when he argues

the challenge of describing determinism to subjects unfamiliar with the concept is daunting to say the least. The description must: (1) make the determinism sufficiently salient, but (2) not trigger fatalistic interpretations, or (3) beg any questions about how to interpret words like “can” and “possibility” and terms like “had to happen”. It is at least arguable that providing an unbiased non-technical description of determinism in half a page is an impossible task.⁹⁵

Yes, the experimental philosopher who wants to study free will is in quite a predicament.

The MPEP, in its varied forms and degrees of severity, is real and worth appreciating. Unfortunately, the methodological problem isn't solved by demonstrating the methodological reliability and validity of a single study or even a group of studies, otherwise we might overcome it without saying much about it—we could just get to work on a few targeted studies and be done with it. Rather, it is a problem that needs to be addressed on a case-by-case basis, and these cases are small in scope. That is, the problem can't be resolved at the level of 'the free will debate,' but more likely at the level of *the debate as to whether laypeople have compatibilist judgments/intuitions when presented thought experiments that refer to singular, mundane actions*. Of course, the methodological problem as a general problem might be considered less problematic if study after study is shown to produce data reflective of people's judgments/intuitions, but

currently no such justification is present for even one or two studies as far as I know. By all outward appearances, the MPEP isn't a call to perform a few more supplementary studies here and there in the realm of x-phi, but to fundamentally reconsider how we prompt, capture, and analyze laypeople's judgments/intuitions.

At the outset of this section, I said I had chosen examples in the x-phi of free will and moral responsibility to suggest ubiquity of the MPEP across multiple areas of philosophical inquiry—and this is the case. Unfortunately, though, the ubiquity of the problem doesn't point to a singular solution at the level of subject matter. That is, showing that this or that study on free will can overcome the MPEP says nothing about other topics given the nuanced differences between philosophical questions and the thought experiments used to get at them. The solution, I will suggest throughout this project, might be resolved by a singular methodological approach, though, one of adding qualitative methodologies to the use of quantitative surveys. It will take that level of coordination to unearth the richness in people's judgments on ethics and beyond.

CHAPTER THREE

A Fine-Grained Look at the Methodological Problem: The “Ought Implies Can” Study Motivating My Work

I. An Overview of the OIC Study I’ve Chosen to Investigate

In this section, I examine a representative x-phi experiment in order to: 1) take a fine-grained look at the methodological problem, 2) introduce the specific “ought implies can” (OIC) study I based my own study on (which is detailed later), and 3) point to the precise features of this OIC study that I believe open it up to the MPEP so that I can motivate the specific design of my own study. All of this helps further my goal of examining which methods best prepare philosophers for understanding the complex judgments laypeople have toward different ethical issues. For my analysis, I will refer to what I call “the ideal story of data collection,” which essentially refers to what *ideally* takes place in an x-phi study that relies on quantitative surveys. In this chapter, my fine-grained exploration involves tracing the ideal story of data collection as it connects with the OIC study; this will help to show precisely which parts are plausibly suspect in this particular study as well as reveal challenges other x-phi studies likely face.

The study to which I refer is by Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard (from here, Chituc et al.), described in their paper entitled “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” and it details three related experiments.⁹⁶ I have chosen the third experiment (which I will refer to as OIC experiment 3) from this study (which I refer to as the OIC study) for much of my focus here, but I am also looking at the study (which includes the three experiments together) more generally.

I chose this research for the foundation of my own experiment and also for this outlining of the ideal story for a number of reasons. First, it functions as a representative study because it includes a standard experimental procedure found in many x-phi studies. Second, it benefits from originating in a study that employs and analyzes qualitative data: while OIC experiment 3 itself *does not* contain a qualitative component, it follows and builds off of the first two experiments in the OIC study, both of which *do* include a qualitative component of special interest to my research. Specifically, in the first two experiments, the researchers recorded brief written statements from subjects by asking them to explain their judgments in a free-write section. While OIC experiment 3 relies solely on quantitative research methods, it can be understood in conjunction with the qualitative data—i.e., the free-write statements—included in the first two experiments within the study, as all three experiments in the study deal with the relationship between the concepts of “ought” and “can.”⁹⁷ This qualitative component is not found in many similar x-phi studies but serves in the OIC study as an informative addition worthy of attention that does not prevent me from speaking about the core of the study, or OIC experiment 3 in particular, as representative of similar x-phi studies.⁹⁸ Third, the OIC study was the focal point of a New York Times piece written by Chituc and Henne entitled “The Data Against Kant,”⁹⁹ which indicates that it has a significant place within public conversation and understanding about what x-phi is and how it functions. This public visibility makes the OIC study a good candidate for analysis as it engages with a conversation that lives both inside and outside of academia.

Lastly, I chose OIC experiment 3 in particular, and the OIC study more generally, because it centers on realistic thought experiments. Here is an example of one of the vignettes from the study, complete with two different endings, which describes a person failing to keep a promise due to plausible circumstances:

Brown is excited about a new movie that is playing at the cinema across town. He hasn't had a chance to see it, but the latest showing is at 6 o'clock that evening. Brown's friend, Adams, asks Brown to see the movie with him, and Brown promises to meet Adams there. It takes Brown fifteen minutes to drive to the cinema, park, purchase a ticket, and enter the movie. It would take 30 min if Brown decided to ride his bike. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time.

[As Brown gets ready to leave at 5:45, he decides he really doesn't want to see the movie after all. He passes the time for five minutes, so that he will be unable to make it to the cinema on time. Because Brown decides to wait, Brown can't make it to the movie by 6. (High Blameworthiness)]

[At 5:30, Brown thinks about riding his bike, but decides it is too cold. Instead, he leaves at 5:45, but his car breaks down five minutes later. He can't fix it himself in time to make it to the cinema, and it is too late to make it by bike. Because his car is not working at the time, Brown can't meet his friend Adams at the movie by 6. (Low Blameworthiness)]¹⁰⁰

Adams' actions are realistic and they take place in a realistic setting; he is not, for example, failing to meet a friend on Mars because he waits too long before stepping into his teleportation device. It is certainly possible that the realism or unrealism embedded in a given thought experiment has little to no effect on a person's judgment regarding some hypothetical scenario. However, it is plausible that people are more often distracted by unrealism, being a departure from everyday circumstances, which causes them to form judgments in response to unrealistic features of a

thought experiment that have no bearing on the philosophical issue being examined. Researchers can avoid such potential distractions by opting for realistic scenarios; the realism, therefore, is something I appreciated in the OIC study.

In their three experiments, Chituc et al. set out to see if people affirm the commonly asserted philosophical claim linked to Immanuel Kant¹⁰¹ that “ought” implies “can,” or that any given person who meets the criteria for a specific obligation (an “ought”) must be able to carry out said obligation (the “can”).^{102, 103} With results in hand, the researchers conclude that Kant’s claim, endorsed by many analytic philosophers, is challenged by the fact that many people appear to judge that “ought” does not necessarily imply “can.”¹⁰⁴ In this way, the collection of people’s judgments is viewed as informative with respect to a particular philosophical theory, namely the one that relates “ought” and “can” in this way.

Like many experimental philosophers, Chituc et al. got their data by administering quantitative surveys with two main parts. The first part is some kind of vignette that briefly describes a hypothetical scenario that is philosophically interesting. In experiment 3, Chituc et al. presented subjects with vignettes describing various situations in which a hypothetical person, Brown, has made some kind of plan either to meet another invented character, Adams, for a movie or to go to the movies alone.¹⁰⁵ In some of these vignettes, Brown fails to meet his obligation with Adams due to circumstances beyond his immediate control—e.g., his car breaks down in transit—while in others he fails due to circumstances that are at first fully within his control—e.g., he decides to wait too long before leaving.¹⁰⁶ Following vignettes such as these, subjects in x-phi studies are presented with the second part of the study: one or more follow-up questions regarding the content in the vignette. In such studies, these questions require quantitatively limited responses from subjects. That is, subjects are restricted to choose from a set of quantitative values, such as

“yes” and “no” or a scaled response, rather than having the option to describe their responses qualitatively in written or oral form. Subjects in Chituc experiment 3 answered questions on a quantitative 100-point scale (on which -50 and 50 mark the extremes and 0 a position of neutrality) in reference to the content in the vignettes they read.¹⁰⁷ In OIC experiment 3, subjects were asked to mark to what degree they agreed with statements asserting that Brown can make his appointment, Brown is to blame for not making the appointment, and Brown ought to make his appointment.¹⁰⁸ Based on responses to such quantitative surveys, experimental philosophers such as Chituc et al. make conclusions about what and how people think about the philosophical issue that is intended to be baked into the presented vignettes and/or follow-up questions (in this case, the veracity of ought implying can).

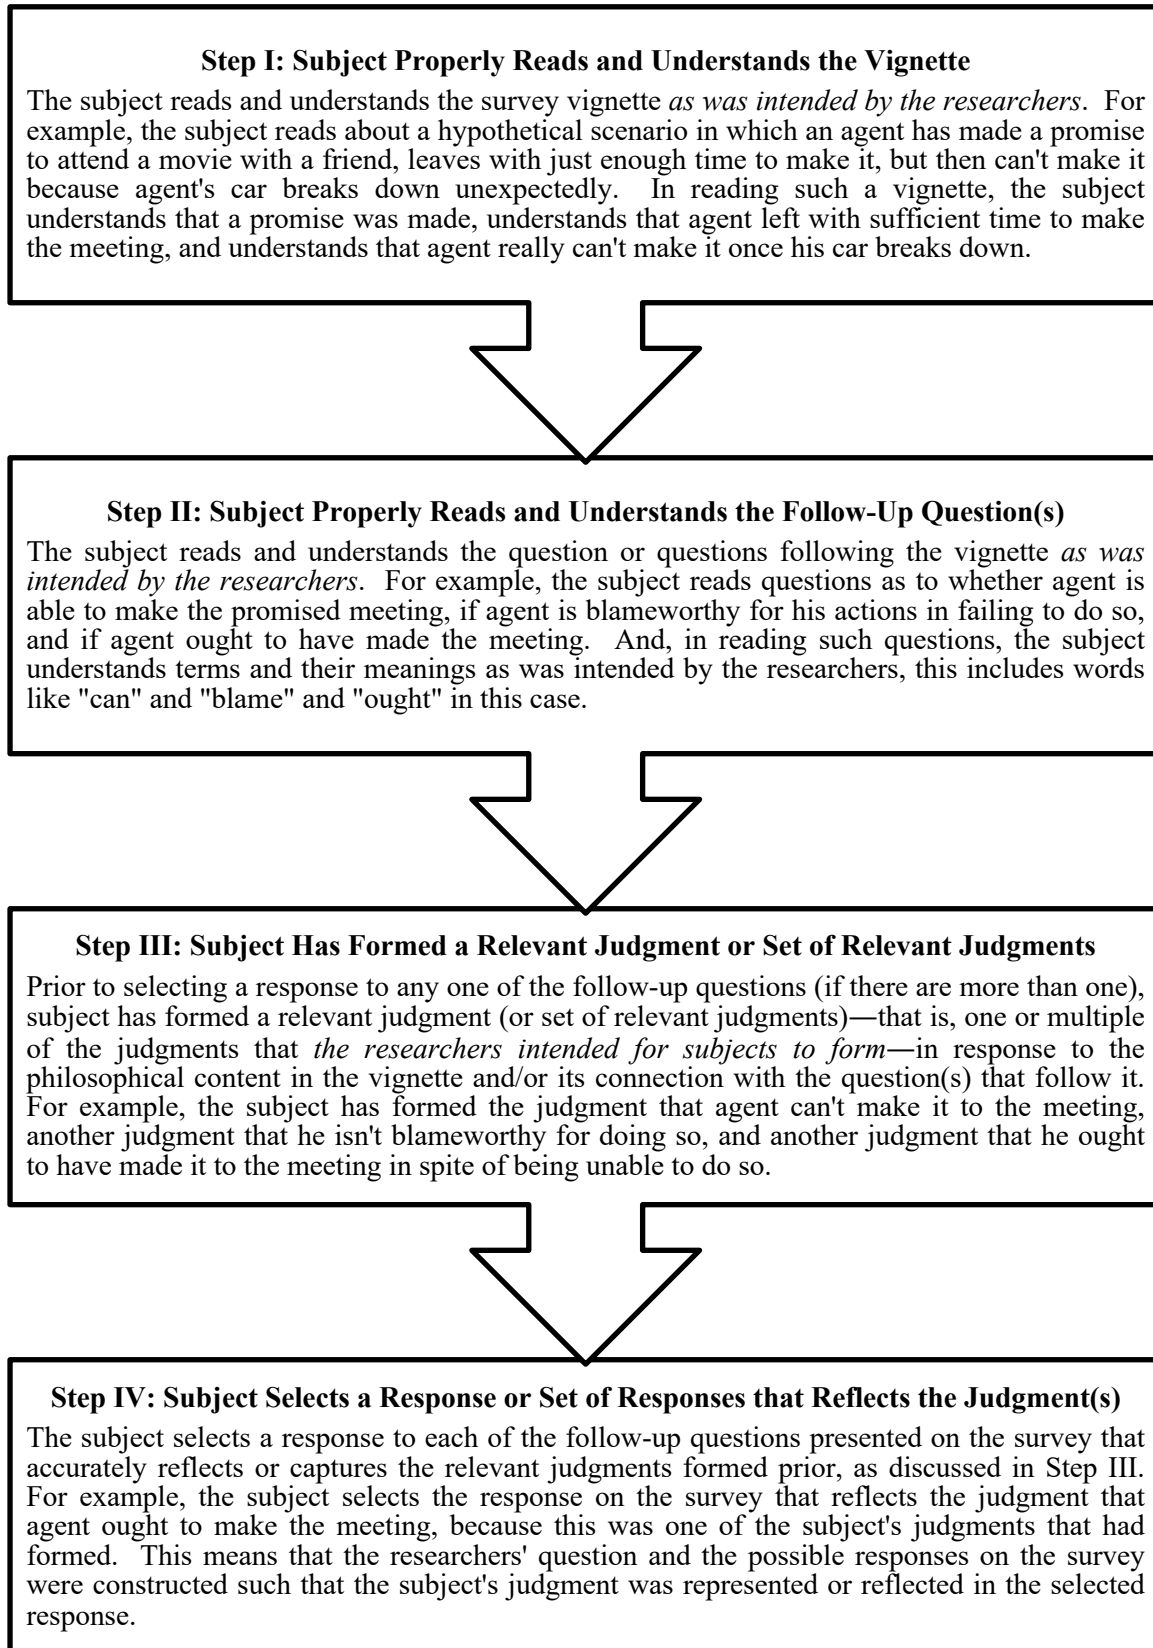
In addition to these two main parts, some experimental philosophers include a qualitative research component in their experimental design which can reveal useful information about subjects’ quantitative responses. For example, in the first two experiments of the OIC study (but not experiment 3), subjects provided qualitative data in the form of open-ended textual responses to the quantitative survey they had completed.¹⁰⁹ Specifically, Chituc et al. asked subjects to explain their judgments that were captured by the quantitative responses they selected in response to the vignettes they read.¹¹⁰ (The responses provided in these open-ended responses do not take center stage for this study, and the reader of the publication is told little of significance that was derived from these responses. This is noteworthy because on the face of it, such an additional qualitative component might appear to serve a similar function as the qualitative methods in my own study, the kind I call for throughout this project, but upon closer investigation it is clear that the researchers did not emphasize this qualitative data nor construct their study around what it might reveal about their quantitative research methodology. In fact, OIC experiment 3, which is

the most robust and meaningful of the three experiments, does not include this qualitative component for whatever reason.) Overall, though, experimental philosophers rarely employ qualitative components such as this one (more on this in next chapter); their primary focus is to capture people's judgments/intuitions via quantitative surveys.

II. X-Phi Experiments Like the OIC Study Rely on an Ideal Story that is Specious

Having outlined the main parts of the OIC study (and thereby studies like it), I now want to discuss the ideal story that links these parts together. I refer to this data collection story as the “ideal story” because I will question its realization. The ideal story, in short, is what x-phi researchers *appear* to assume happens, beat for beat, during an experiment. If studies such as the one by Chituc et al. stated explicitly and precisely that *this* or *that* is what is thought to occur during data collection, then I would not have to claim that this research story is grounded on what *appears* to be an assumption (rather I could cite the authors stating precisely what assumptions they are and are not making), but no such disclaimer precedes this study.¹¹¹ Figure 1 below details the ideal story of data collection found in the kinds of x-phi studies I am focused on, which are standard in the field of experimental philosophy. Embedded in the descriptions of Figure 1 are examples that apply specifically to OIC experiment 3, the one that serves as the basis for my own study (which I detail later).

Figure 1: The Ideal Story of X-Phi Data Collection¹¹²



What is uncontroversial about this story is that subjects provide responses on the surveys they are given (as detailed in Step IV). In addition, these responses are *in response* to the vignettes and the follow-up question(s) detailed in Steps I and II: There is no reason to think most subjects choose their responses capriciously, or that they base them on thought processes that are completely divorced from the thought experiments they read. However, it is unclear whether the rest of this story maps onto actual events in the x-phi “lab”—which can refer to the physical space where an experiment took place or to the online space such as Amazon Mechanical Turk. Specifically, the questions listed below (*i* through *viii*) remain largely unanswered, or at least unsatisfactorily answered. This is true *even when* we assume that for any given question every part of the ideal story other than the one being questioned holds true. Each of the questions below could be answered by empirical data—like the data I collected in this present dissertation with my own study—that speak to an actual story quite different than the ideal one. I do not invoke these questions for the purpose of denying the ideal story abstractly, but for the hope that the story can be corroborated definitively. These are the kinds of questions, then, that I sought to address with my own study (which I will detail later). Again, before discussing the results of my research, it is important to note the openness of these questions so that my results aren’t taken as obvious once viewed with the benefit of hindsight. I justify each question below with my rationale for bringing it up:

- i.* **Question:** Is a subject’s judgment (that has formed *by* Step III) being adequately captured by his or her response(s) in Step IV given the limited responses available on the follow-up question(s)?
Justification: It is plausible that participants’ judgments—being complex, nuanced, grand, etc.—are being pushed through a limited number of choices *or* through a variety of choices that themselves restrict or adulterate a given judgment.
- ii.* **Question:** Would a subject choose a different response in Step IV if there were more or other responses available?

- Justification:** It is plausible that a subject’s response in Step IV would be different if he or she were presented with more or different choices in response to a given follow-up question.
- iii.* **Question:** Are subjects’ responses in Step IV being undesirably shaped by subjects’ reading of the follow-up question(s) or the subjects’ reading of the available responses to the follow-up question(s)?
Justification: It is plausible that some subjects form a judgment based on the vignette (prior to the follow-up question reading) that is shaped or modified by the reading of the questions themselves or the choices provided, and thus a subject’s response to the follow-up question(s) in Step IV reflects a now-modified judgment that emerged after the reading of the follow-up question(s) or after the reading of the possible responses rather than from some initial impression of the hypothetical case itself. That is, had similar questions with different wordings or with different available responses followed the vignette description of the hypothetical case, then subjects might have chosen different responses in Step IV.
- iv.* **Question:** It is unclear what exactly is thought to *cause* a subject to form a judgment: the vignette, the follow-up question(s), the available responses to the follow-up question(s), some combination of the three, or something else. If researchers do not know, then how can they make meaningful claims about the connection between a subject’s judgment and the material he or she is presented with?
Justification: It is plausible that subjects form one or more judgments from any one of the parts of the material they are supplied with in a study, and therefore it is plausible that a judgment formed based solely on the vignette is changed or negated by a judgment that formed based on the subsequent reading of the first follow-up question which might be changed or negated by a judgment that formed on the subsequent reading of the possible responses to the first follow-up question.
- v.* **Question:** Are subjects’ cognitive responses to the survey experience, those that are labeled “judgments” in Step III, sufficiently uniform such that they can be accurately understood as “judgments” which can be compared to other “judgments” from other subjects within the study or those from other people outside the experiment?
Justification: It is plausible that subjects’ cognitive responses to the content relevant to the philosophical idea being studied are varied and share few relevant qualities such that it would be inaccurate to refer to them and analyze them as “judgments.”
- vi.* **Question:** Is a subject’s response to a follow-up question in Step IV *caused* by the judgment formed prior to Step III?
Justification: It is plausible that subjects’ responses to the follow-up question(s) are caused by a whole range of things other than the relevant judgment such as a desire to select what is perceived to be the ‘right’ response given the context of the experiment.
- vii.* **Question:** Is a subject’s judgment that formed prior to Step III in direct response to an implicit or explicit understanding of the philosophical idea that is designed to be

captured by the relevant content in the vignette and the follow-up question(s) read during Steps I and II?

Justification: It is plausible that subjects are reacting to a misunderstanding of the philosophical idea that is designed to be captured by the relevant content even though their judgments are in direct response to the relevant content in some misinformed way.

viii. Question: Are the vignettes and follow-up question(s) effectively causing the subjects to form a judgment in response to the relevant philosophical idea described therein?

Justification: It is plausible that subjects form judgments in response to features of the vignettes and follow-up question(s) other than the relevant philosophical idea, such as a perceived notion about what is “really being asked” or a reaction to unimportant features of the hypothetical case.

In addition to these questions, there are likely many more that can help researchers better understand whether x-phi studies are achieving what they intend to achieve.

It is possible that experiments like those in the OIC study successfully produce data in line with the ideal story, but it is plausible they do not. The questions above point to a number of uncertainties and vulnerabilities in this story that warrant further investigation. Chituc et al. *may* be capturing people’s judgments about whether “ought” implies “can,” but they may instead be capturing a variety of thought processes only vaguely related to the scenarios subjects are presented with, thought processes problematically funneled through the limited survey responses made available. And because these same worries apply to all x-phi studies that rely on these quantitative survey methods, I thought it would be beneficial to study this data collection process closely to either confirm the veracity of the assumptions that underlie it or expose them as unfounded. And so my own study, detailed later, is designed to get a better look at this ideal story, in hopes of recommending which experimental methods work best and in hopes of better understanding how people think about this foundational ethical principle. In the next section, though, we can begin this better look by scrutinizing the small amount of qualitative data reported in the OIC study. I will show that even there we see evidence of the ideal story breaking down.

III. A Systematic Analysis of the OIC Study Shows that it Fails to Follow the Ideal Story

Before getting to my own study on OIC experiment 3, there is still more that can be unearthed by digging into the details of the OIC study and its data. In the end, this section should stand as the most clear and compelling evidence of the MPEP (prior to my study) since it is the one in which we can look most closely and analytically at x-phi data.

In experiment 1 of the OIC study, each participant read the same initial sentences: “Adams promises to meet his friend Brown for lunch at noon today. It takes Adams thirty minutes to drive from his house to the place where they plan to eat lunch together.”¹¹³ In the low blame condition, participants subsequently read the following: “Adams leaves his house at eleven thirty. However, fifteen minutes after leaving, Adams [*sic*] car breaks down unexpectedly. Because his car is not working at that time, Adams cannot meet his friend Brown at noon, as he promised.”¹¹⁴ For subjects in the high blame condition, they read the following paragraph just after the initial sentences that all participants read: “Adams decides that he does not want to have lunch with Brown after all, so he stays at his house until eleven forty-five. Because of where he is at that time, Adams cannot meet his friend Brown at noon, as he promised.”¹¹⁵ Following the respective vignettes, participants in the low and high blame conditions were asked to what degree that agreed or disagreed with the following statement: “At eleven forty-five, it is still true that Adams ought to meet Brown at noon.”¹¹⁶ And, following, participants were given the free-write space to explain their responses (more on this below).

The core finding was that “participants were more likely to say that an agent ought to keep a promise they can’t keep in the *high blame* condition...than in the *low blame* condition.”¹¹⁷ They argue this supports their hypothesis 2 (H2) that “participants will judge that an agent ought to do something that the agent can’t do when the agent is to blame for the inability” over their hypothesis

1 (H1) that “participants will deny that an agent ought to do something that the agent can’t do, regardless of whether the agent is to blame for the inability.”¹¹⁸ In short, they argue their data support the hypothesis that “ought implies can” isn’t an iron-clad principle people uphold since they appear to attenuate their support of it when an agent does something blameworthy to bring about their failure to be able to act on an obligation. And, to emphasize this finding, Chituc et al. make reference here to the qualitative data in brief. Specifically, they claim that “some of the participants outright rejected H1 in their explanations: e.g., ‘Brown is still going to be waiting for him at noon. Adams won’t be ABLE to but he still OUGHT to’ (capitals in participant response).”¹¹⁹ This is the only free-response example they quote directly and it is one that supports their claim that participants explicitly denied H1, but as we look into other free-response examples, we will see that other participants gave responses that aren’t so straightforwardly in support of H1 or H2.

Without making too much a fuss here, it is a bit frustrating that this one free-response is singled out since it clearly supports the hypothesis the quantitative data appear to support, and no mention is made of a free-response like this one which appears to challenge the trend toward H2: “he can not [*sic*] drive 30 minutes in 15 minutes.”¹²⁰ This participant, who importantly read the high blame condition first such that one can’t worry she wrote this after carrying over some sense of OIC from the low blame condition, appears to be supporting something very much like H1 with this response viewed in combination with her selection of “-42” on the scale between negative 50 and 50 as a measure of her agreement or disagreement with the statement that “at eleven forty-five, it is still true that Adams ought to meet Brown at noon.”¹²¹ If qualitative responses in support of one hypothesis are highlighted and quoted in an analysis of data, those in opposition to it be highlighted as well if the researchers are wanting to claim that one hypothesis is better supported.

Now, moving closer to my analysis of the qualitative data in experiment 1 of the OIC study, there is a bit more to summarize. After citing the one example of qualitative data that aligns with their quantitative data, Chituc et al. do interestingly note that some free-responses included sentiments that lead them to revise their vignette for the next experiment, thereby implicitly affirming some kind of value in the qualitative data for revealing something like the MPEP. Specifically, they note “some participants provided alternative actions that Adams should have done instead, such as calling his friend,” which problematically suggests some participants might have been conveying that “Adams should still meet his friend, even if he’s late.”¹²² Chituc et al. seem to accept, based on certain qualitative responses, that they didn’t have sufficient epistemic justification to conclude that their quantitative data really did capture the relevant philosophical judgments needed to support a claim about what participants did or didn’t judge to be the relationship between “ought” and “can” in relation to this particular case. Specifically, it seems that some participants were not tracking the plot of the vignette in an important way: they didn’t seem to be responding to the question of whether Adams ought to make his obligation at noon but rather they responded to the question of whether he should do something like the next best thing. And, if this is in fact how some participants interpreted the scenario, then it seems clear that their judgments speak to *what someone ought to do in a situation wherein an obligation can no longer be met as was planned* (which *doesn’t* pit “ought” against “can” in the right way), rather than to *whether or not someone ought to do something he or she can’t do* (which *does* pit “ought” and “can” together so as to challenge the Kantian principle of interest).

While Chituc et al. don’t specify which free-responses prompted their vignette revision, it is presumably examples such as the following: “He should call Brown and ask him to pick him up or to wait until he can get his car running.”¹²³ This participant offered an 8 on the scale between -

50 to 50 in support of the claim that “at eleven forty-five, it is still true that Adams ought to meet Brown at noon.”¹²⁴ In other words, this participant responds quantitatively that Adams ought to do something that, from the vignette, he is no longer able to do, and yet, she qualitatively expresses something that makes this affirmation of obligation seem like an affirmation that Adams should take other actions given that making it on time is no longer a possibility.

And while it is encouraging that Chituc et al. take the qualitative data as indicative of something *beyond* the quantitative data, I think it important to highlight that their conclusion from free-responses like this one—assuming that this is in fact one of the ones that caught their attention—is that proponents of the Kantian principle—as if they would be the only ones to object to this data!—might argue that such responses indicate that participant *misunderstood* the vignette—or reacted differently than the researchers intended—and responded as if they were being asked to say if Adams ought to take action, any action, to best make it to his appointment as close to the original meeting time as possible or ameliorate the situation as best as possible or something similar.¹²⁵ What they don’t acknowledge, though, is that this participant’s responses (the mentioned quantitative one plus its corresponding free-response), and others like it, might actually *evidence the opposite* of what Chituc et al. took the total quantitative data to show: these such responses might not be revelatory of a vignette in need of revision, but instead revelatory of the fact that some participants, given that they search for alternative actions for Adams, are expressing, as best as their limited available responses allow them to, that “ought” really does imply “can” or something similar.

Let me elaborate: this participant’s responses (the mentioned quantitative plus the corresponding free-response) might actually show that this participant *did* in fact grapple with a relevant philosophical judgment. Specifically, she perhaps thought it impossible for Adams to

make it to the meeting, so much so that she came up with alternative explanations *given* that impossibility. So, while Chituc et al. might want to view this occasion as invalid or erroneous or based on a misunderstanding, it might be that this participant's responses were reflective of a genuine philosophical judgment regarding OIC, namely that "ought" *does* imply "can" such that when presented with a case where one is being asked if someone ought to uphold an impossible obligation, she concludes that the only meaningful response would be to interpret "ought" as asking about other possibilities, which would be a very different sense of "ought" than is embedded in the Kantian principle. Whereas questions of the Kantian OIC are grounded on whether or not someone ought to do something when he or she can't, the sense that participant might have been tapping into is one in which the question is *what* ought someone do, not *whether or not* he or she ought to do something. That is, I am speaking of the difference between 'Ought he make it?' and 'What ought he do?' Of course, experimental philosophers want their data to reflect a *clear* distinction between this or that kind of judgment, and they shouldn't have to wonder if a response that appears one way at first might reflect the opposite judgment upon closer examination. I am not, therefore, arguing that this vignette needs no revision given my interpretation that the participant may have actually formed and responded based on a genuine philosophical judgment. Unfortunately, though, this interpretation points out just how difficult it could be to interpret a given participant's response and how to revise a vignette so as only to get the kinds of judgments one wants reflected in quantitative data alone. This brings us back to the problem of relying solely on quantitative data in x-phi projects like this one by Chituc et al.: I just don't know what thought experiment you could design such that every (or most every) quantitative response it produced would be clearly reflective of a certain philosophical judgment *unless* it was probing for judgments regarding the most straightforward cases (recall Paul the murderer) which make up few if any of

the philosophy debates for which we might be interested in people’s potentially varied philosophical judgments. Chituc et al., in their second experiment in the OIC study, do make revisions in hopes of producing data *more* reflective of the sought-after kinds of judgments, but we will again see (below) how new problems—insurmountable if one just wants to keep throwing quantitative surveys at them—arise that point again to the MPEP.

I want to finish my analysis of experiment 1 in the OIC study by looking at more examples of free-responses so as to 1) avoid the sense that I singled out the only interesting free-responses, and 2) further demonstrate that there is enough in these free-responses for us to see the value in collecting more qualitative data, beyond a short free-response, in similar x-phi experiments. Below, I have listed a number of qualitative responses from the first experiment that I found noteworthy and a brief explanation of what I find noteworthy in each. While I am pointing out a number of problems indicative of the MPEP and the breakdown of the ideal story I also want to remind the reader that some of these problems appear to be precisely why Chituc et al. made changes and completed the second experiment with new considerations in place. That is, it is not as though Chituc et al. would necessarily find all my pointing to uncertainty in the data disagreeable. Still, what is odd is that Chituc et al. seem to take qualitative responses seriously enough to revise on their behalf, but not so seriously that they feel that they might be fundamentally unjustified in taking their quantitative data at face value.

1. Experiment 1 Participant #43, Low Blame

a. Quantitative Response: 50

b. Qualitative Response: “There’s a difference between ‘ought to’ and ‘can.’

Plus, ‘ought to’ can have several different connotations. I read it less like

he's going to and more like that's his goal, even if now his goal can't be met."¹²⁶

- c. **My interpretation:** While it is possible that this subject is one of few who experienced some lack of clarity as to what “ought” means in the survey, it is likely that many subjects selected their responses based on *one* connotation of “ought” over another. And, depending on the connotation, a subject might form and report very different judgments about the relationship between “ought” and “can.” Chituc et al. are aware that connotation might be a problem for their conclusions, yet they dismiss it as a significant issue for their present study. When discussing their results, Chituc et al. write that “there is no reason to posit different meanings of “ought” other than to preserve the truth of ‘ought’ implies ‘can’ in light of contrary evidence.”¹²⁷I disagree. This participant’s comment alone serves as reason to question whether subjects are responding to some intended meaning of “ought.” In fact, it strikes me that it would be responsible in this situation for Chituc et al. to acknowledge that at least *one* subject articulated that she was completing the survey based on *an* understanding of “ought” while acknowledging that others exist. Even if their experimental design did not allow for them to *do* anything about this source of ambiguity, it would be responsible for them to note it as a pitfall.

2. Experiment 1 Participant #16, Low Blame

- a. **Quantitative Response:** 24

- b. **Qualitative Response:** “There's still time and he has a legitimate reason for being a little late.”¹²⁸
- c. **My interpretation:** This one is noteworthy because participant seems to suggest that Adams *can* make it on time—which he *can't* according to the vignette—and that it is understandable that he might be late—which seems to show participant thinks Adams will *not* make it on time. It seems to me that the ideal story breaks down at some point since participant doesn't appear to be reacting to the *impossibility* of Adams' making it on time, so it is difficult to see how the judgment that participant forms is connected to the philosophical content of the thought experiment, namely the relationship between “ought” and “can.”

3. Experiment 1 Participant #30, Low Blame

- a. **Quantitative Response:** 50
- b. **Qualitative Response:** “It is still true because he made a commitment and just because he would not get there in time it doesn't change that”¹²⁹
- c. **My interpretation:** I want to grant to Chituc et al. that this response shows a clear-cut and robust endorsement of the notion that “ought” doesn't imply “can” since participant says the commitment remains despite Adams' inability to make it on time. But I can't without knowing more. Specifically, I'm worried that participant's language that “[Adams] would not get there in time”¹³⁰ is participant expressing that Adams still has a commitment to maintain, one of getting to the meeting *even if it means getting there late*, and that participant isn't saying that Adams ought to do

what he can't do. Still, this gets about as close as we can get to a rejection of OIC in a participant response of this kind.

4. Experiment 1 Participant #57, Low Blame

- a. **Quantitative Response:** -12
- b. **Qualitative Response:** "He is physically prevented from meeting his friend at noon, so he no longer 'ought' to"¹³¹
- c. **My interpretation:** I just expressed much skepticism toward the previous participant's response that seemed to get quite close to rejecting OIC, but I think this one pretty squarely shows endorsement of OIC. In part, that's because OIC is easier to endorse than reject in some sense, especially when one makes it clear that he or she recognizes a physical impossibility preventing the "can" part of the principle. I still would want to find more out from participant, but this seems like an OIC endorsement in that Adams is explained to be no longer under obligation of an "ought" for just the reason that physical impossibility makes the "ought" no longer obtain. If the previous participant had expressed a similar focus on physical impossibility, know that I would easily grant that it is an anti-OIC response just as I grant a pro-OIC response in this example. For example, had the previous participant said that *despite being physically unable to make it on time Adams still ought to*, then I'd be happy to say we have an anti-OIC judgment.

5. Experiment 1 Participant #8, High Blame

- a. **Quantitative Response:** 26

- b. **Qualitative Response:** “He ought to meet Brown at noon, however it is impossible for him to do so because the earliest he could arrive is at 12:15. He had made a promise to Brown and should have followed through with it.”¹³²
- c. **My interpretation:** I really do want to grant that this judgment is anti-OIC. Participant clearly articulates that Adams can’t make it on time and still endorses quantitatively that he thinks Adams ought to make it. But then, participant *clearly states* that Adams “should have followed through with it.”¹³³ That is, it’s unclear if participant is saying that at 11:45 Adams ought to do things that he can’t—which seems to fall out of the quantitative response selected—or if in general and in total Adams *should have* met his obligations, which in a general sense *he was able to do* had he not decided on waiting around. That is, we can say, Adams *was* able to make his obligation, but then decided to wait too long, and participant might be expressing that Adams *should have* taken actions to meet the obligation and that *he should not have* waited around. Since the questionnaire doesn’t leave room for a response of this kind, we are stuck with a numerical response to a claim about 11:45 in hopes that participant wasn’t funneling a more complicated response through the only available question/scale. That is, there is no quantitative response equivalent to *Adams ought to have made it because he promised he would, but at 11:45 he no longer can so he no longer ought to.*

6. Experiment 1 Participant #19, High Blame

- a. **Quantitative Response:** -37
- b. **Qualitative Response:** “He doesn't want to”¹³⁴
- c. **My interpretation:** This one is especially noteworthy because it brings up a possible interpretation of Adams' scenario that might just bypass what we philosophers are interested in altogether: Adams doesn't maintain an “ought” because he clearly doesn't want to make his meeting. By waiting around, Adams makes it clear that he doesn't want to meet up, so obligation really isn't his problem anymore. Should Adams go? No, he clearly doesn't want to. There is no tension here between “ought” and “can” because the “ought” falls away now that it becomes clear Adams has no desire to meet. Simply put, one should not do what one does not want to do, so if participant interprets the question as asking *if* Adams ought (should) go, then the answer is clear: no, because he has no desire or intention to. Ought I buy a coffee? No, because I don't want coffee.

7. Experiment 1 Participant #39, High Blame

- a. **Quantitative Response:** 38
- b. **Qualitative Response:** “Adams made the decision to neglect his obligation. He should be held accountable.”¹³⁵
- c. **My interpretation:** This is noteworthy because it seems like participant might be agreeing—with a 38—that Adams has an “ought” because participant takes “ought” to express that Adams is “accountable” for his actions.¹³⁶ This strikes me as a different sense of “ought” than the one in

the Kantian principle wherein “ought” isn’t identical to accountability or blameworthiness. If participant is expressing a kind of blame here, then this response is hardly something that could be used to support a negation of OIC. In other words, I think this is further evidence that participants might be understanding “ought” in different ways, as this participant clearly aligns “ought” with accountability, maybe even a general sense of accountability that blankets over Adams’ total actions, which is different than “ought” as a moral obligation that Adams may or may not be held blameworthy or accountable for.

8. Experiment 1 Participant #43, High Blame

- a. Quantitative Response:** 50
- b. Qualitative Response:** “In this context, I read ‘ought to’ as meaning it's what he should do, it's what he promised to do. Even though he doesn't want to.”¹³⁷
- c. My interpretation:** And lastly, this one is noteworthy because participant echoes the sentiment that “ought” might read different ways by the fact that she claims to have understood it in one particular way. In addition, her response seems to comment on what Adams should generally do, because he made a promise, not what he *ought to do* precisely at 11:45 when he could no longer do it. It is just difficult to conclude that participant was fully appreciating the impossibility at 11:45 when she comments on the promise and his lack of desire to fulfill it rather than the impossibility of his predicament at 11:45.

I just dug into these qualitative responses, short as they are, and I think I've successfully shown the MPEP is alive and well, that the ideal story is specious, and that qualitative data in small amounts say a lot and therefore in abundance likely say even more. To avoid belaboring the points I've made, I'll just cut through as quickly as possible to some interesting qualitative responses in the second of the OIC experiments after briefly outlining its general quantitative findings. Then, we'll be on our way to look at OIC experiment 3 in great detail in chapter five when I discuss my partial re-creation of it complete with much qualitative data.

To further test the relationship between blame, ability, and obligation, Chituc et al. carried out a second experiment as part of the OIC study. In large part, the thinking behind the second experiment is in response to a potential objection to the results from the first, namely that participants' judgments regarding 'ought' were influenced by "*blame validation*" because such judgments were "distorted by a motivation to hold wrongdoers accountable in the *high blame* condition."¹³⁸ This objection really has two thrusts: 1) high blameworthy actions bias people to make judgments that endorse "ought"—even when someone lacks ability—so as to hold the wrongdoer accountable, and 2) less blameworthy actions aren't so biased by this desire for accountability such that a wrongdoer in the low blame condition is being viewed without any biasing factors and therefore what participants report regarding the low blame case is an accurate representation of their views on the relationship between "ought" and "can." This objection actually maps onto my analysis of participant #39 above, who appears to be viewing "ought" in just this way: a form of accountability.¹³⁹ And, according to Chituc et al., this objection suggests an empirically testable claim, namely "if judgments of *ability* and *blame* are collected as well as *ought* judgments, then judgments of *ought* and *ability* would show a relationship in the *low blame* condition."¹⁴⁰ In a footnote explaining a possible issue with variance given the scale-nature of the

survey questions, Chituc et al. elaborate on this relationship by conveying that the truth of OIC would have participants with stronger judgments of inability also have stronger judgments of obligation.¹⁴¹ But, they go on, if H2 is correct, then there is a different result we should see, namely that regardless of the blame condition we should expect participants' agreements and disagreements with whether Adams "ought" do something to track the blameworthiness of his actions rather than whether or not he is able to carry out his obligation.¹⁴²

We might have issues with the framing of the objection in support of H1 or the prediction in support of H2, but nonetheless this is how Chituc et al. see matters. So, they construct a single vignette modified from experiment 1 in hopes of making it clear to participants that Adams has no real solutions at his disposal once his plans go awry:

Brown is a CEO of a large company in the economic boom in the middle of the 20th Century. At 2 o'clock, Brown has a meeting in the city to make a significant financial decision that will decide the future of his company. Since so much money is at stake, he asks his trusted personal advisor, Adams, to meet him on the 12 o'clock train. On the train, he plans to discuss his decision on the ride into the city, where Brown will go straight to his 2 o'clock meeting. Adams promises to meet Brown on the train at noon. It takes Adams thirty minutes to drive to the train station, park, purchase a ticket, and board the train. However, fifteen minutes after leaving at eleven-thirty, Adams [*sic*] car breaks down unexpectedly. Because his car is not working at the time, Adams cannot meet Brown at noon, as promised. Since cell phones have not been invented yet, Adams has no way to contact him.¹⁴³

Subjects were asked three follow-up questions, based on the same 100-point scale from the first experiment, each presented as statements that participants were asked to agree or disagree with:

“(i) at 11:45 AM, Adams ought to keep his promise, (ii) Adams can keep his promise, and (iii) Adams is to blame for not keeping his promise. The questions were presented in random order.”¹⁴⁴ The phoning for help option is removed, the blameworthiness of Adams’ actions are less than when he waited around, and participants are asked to grapple with the actions of a man who leaves 30 minutes ahead of time for a trip that takes precisely 30 minutes.

The authors summarize the quantitative results as supportive of H2 because people’s “ought” judgments tracked “blame,” not “ability” as OIC would predict.¹⁴⁵ At face value, there is nothing objectionable here, but again I want to look at the qualitative data to bring some more complexity and accuracy into this whole picture. Chituc et al. make no mention of any of the qualitative data from this second experiment, which I find highly unfortunate since it is rich with ambiguous and doubt-motivating nuance. I’ve chosen some noteworthy responses that by my lights cast significant doubt on the ideal story holding up for the production of the summarized quantitative data.

1. Experiment 2 Participant #68

a. Quantitative Responses

- **Ought:** 50
- **Can:** -47
- **Blame:** 23

b. Qualitative Response: “Adams is physically unable to get there on time, but ought to be- it's expected of him.”¹⁴⁶

c. My interpretation: This participant seems to be directly responding to the question of whether “ought” implies “can” in a way that I, as a researcher, would feel, at least prima facie, confident that my survey was tapping into

the right kind of judgment. First, I'd note that participant uses "ought" in the response *and* provides a follow up explanation of how this is being understood (so it seems). Secondly, he references the physical impossibility of Adams' actions such that we can quite confidently conclude that participant is leveraging a robust sense of "can't" against this "ought" and still finds the "ought" to obtain for Adams' inability to follow through—here, "ought" doesn't appear to imply "can." Better yet, the participant completes the questions in just the way we'd hope: he disagrees that Adams can make it, agrees that he ought to, and somewhat agrees that he is to blame. But even here, there's still some ambiguity such that even if *all* my participant responses in a study looked this clean, which for Chituc et al. they certainly don't, we could still wonder if subjects were thinking of "ought" in just the ways philosophers do when they discuss OIC. For one, the qualification regarding "ought" might reveal that subject thinks that "ought" has to do with an overarching social obligation rather than anything like the binding "ought" that obtains for a now impossible action. When subject clarifies that "it's expected of [Adams]"¹⁴⁷ we are left to wonder which of the following this translates to: 1) in some general sense, those who make promises are expected to fulfill them; 2) in some general sense, in society we expect people to meet their promised appointments; 3) Adams' friend is expecting him to make it even though he can't actually do so, but were Adams' friend to know that he physically can't make it, he would still express some sense of *you ought to be here* but this would be

more of a way of conveying *you ought to have taken the necessary steps so as to be here and my feeling of blame doesn't dissipate the moment you can't make it even if you can't any longer in a metaphysical sense be obtained by an "ought" such that you ought to uphold your promise*; 4) just in the way OIC understands the issue. I really am not being nit-picky here: OIC is more complicated than just seeing the words *he ought to even if he can't*.

2. Experiment 2 Participant #5

a. Quantitative Responses

- **Ought:** 50
- **Can:** -50
- **Blame:** -50

b. **Qualitative Response:** "If you make a commitment, you 'ought' to fulfill that commitment, but there can be things beyond your control that prevent it, for which you are not to blame."¹⁴⁸

c. **My interpretation:** Quantitatively, we have a perfect expression that "ought" doesn't imply "can" since Adams is said to be fully unable to meet the obligation and still participant says fully that he ought to. But the qualitative response challenges this picture since participant seems to be expressing something general like *one ought to follow through on promises*. This hardly seems like a negation of OIC since participant isn't claiming Adams should literally do something that he can't do, but instead claiming—it seems—that people should uphold their promises. The ideal

story breaks down somewhere here because it doesn't seem that participant is reporting a judgment in response to the philosophical content intended to be embedded in the survey, which is the relationship between "ought" and "can" in the way understood through Kant's OIC.

3. Experiment 2 Participant #38

a. Quantitative Responses

- **Ought:** 37
- **Can:** -37
- **Blame:** -38

b. **Qualitative Response:** "i only answered agree to adams ought to meet brown because he told he him would, i answerd disagree to the other two questions because i really dont think adams could have prevented his fate and its just an unfortunate event."¹⁴⁹

c. **My interpretation:** Similar to the previous participant's response, this one seems to be saying something like *you ought keep your promises*.

4. Experiment 2 Participant #9

a. Quantitative Responses

- **Ought:** -14
- **Can:** -40
- **Blame:** -46

b. **Qualitative Response:** "He's not to blame for not being able to meet Brown at noon, since his car broke down, which was out of his control. I guess he can still meet Brown somewhat, if he manages to find an alternate mode of

transportation. He ought to meet Brown at noon to a certain extent, but since he's unable to, I wouldn't say he ought to, to a large degree.”¹⁵⁰

- c. My interpretation:** What can we make of nuanced responses like this? There are aspects here that seem both supportive and negative of OIC. Plus, we get a qualitative suggestion that Adams might be able to make it in conjunction with a quantitative response disagreeing that he can. I think two things are clear: 1) responses like this can't be treated as straightforward endorsements of certain quantitative values, and 2) while we might hope that few judgments are this nuanced given that few participants express so much in their free-responses, it is entirely possible that such judgments are common even if underreported.

5. Experiment 2 Participant #18

a. Quantitative Responses

- **Ought:** -31
- **Can:** -24
- **Blame:** -19

- b. Qualitative Response:** “Adam's car has broken down. It will be difficult for him to meet up with Brown, unless he can get a ride from passer byer.”¹⁵¹

- c. My interpretation:** This participant's responses are important in subtle ways. In their analysis of this experiment, Chituc et al. say that they ran a statistical analysis additional to the first one they report, after removing data from 29 participants “whose *can* ratings were above the midpoint, suggesting that, contrary to the vignette, these participants believed that

Adams could keep his promise.”¹⁵² They ended up reporting that this didn’t affect their statistical analysis much¹⁵³, but I want to highlight the fact that with this move they show their sensitivity to people’s responses that seemed to get Adams’ inability completely wrong in their quantitative responses. The current participant is noteworthy because he shows that being sensitive only to quantitative data, and not to qualitative data (which Chituc et al. don’t highlight at all in response to this second experiment) can be problematic. This participant disagrees that Adams can make it—with a score of -24—but nonetheless qualitatively reports that there is a chance that Adams can make it. Not only does this free-response negate information that is clear in the vignette—Adams can’t make it, really—it also provides us with no sense of how this ambiguous degree of inability that Adams is experiencing *affects* the “ought” judgment. Without knowing more about this participant’s thinking, we might actually know less about his thinking than were we to know nothing at all since his data might be tempting us into thinking that *we’ve got some idea* of his philosophical judgment toward “ought” and “can” when we don’t.

6. Experiment 2 Participant #52

a. Quantitative Responses

- **Ought:** 28
- **Can:** -50
- **Blame:** -50

b. Qualitative Response: “The scenario states that he left via car at 11:30 and broke down 15 minutes later. Since it takes him 30 minutes to get from home to board the train it will require another 15 minutes of driving which he is now unable to do. No other means is available as inferred from the passage indicating that ‘Adams cannot meet Brown at noon’. Since Adams presumably didn't sabotage his vehicle on purpose Adams is not to blame for missing the noon meeting with Brown. Despite the situation, the original expectation of meeting at noon still exists and is still subject to the expectation especially since Brown is unaware of the car issue. As such, Adams still ‘ought to meet Brown at noon’ with a reasonable degree of expectation via alternate methods of getting to the noon meeting- albeit not at full expectation since there is still a chance that alternate means are not available.”¹⁵⁴

c. My interpretation: This participant offers a detailed explanation of his choices. This is noteworthy beyond this atypical detail because participant appears to be linking “ought” with “expectation” which he links in one instance with what Brown expects given his lack of knowledge regarding Adams’ car troubles. Does the participant philosophically understand “expectation” to operate the same as “ought”? Or might it be that participant is using the question about “ought” as a means to convey his judgment that Adams is expected to be on time, an interpretation that amounts to the breakdown of the ideal story in that a judgment of “expectation” is being funneled through limited questions about “ought,” “can,” and “blame.” I

think it is clear that if a participant means by his or her agreement with the “ought” question that Adams *is expected to arrive on time* then we wouldn’t feel as if we have in our possession a judgment regarding OIC but instead about expectation. If Brown is unaware of Adams’ trouble, then of course Brown still expects Adams’ arrival.

7. Experiment 2 Participant #79

a. Quantitative Responses

- **Ought:** 13
- **Can:** -39
- **Blame:** 32

b. **Qualitative Response:** “Adams is responsible for his own possessions being in working order.”¹⁵⁵

c. **My interpretation:** This one is noteworthy just in that what Chituc et al. consider low blame isn’t necessarily low blame to all participants. While many people certainly hold that car troubles are completely not the fault of the car owner, others disagree. Proper maintenance of a car reduces the probability of a breakdown, and so some participants, like this one it seems, might think themselves justified in blaming Adams for leaving his car in precarious condition when he relies on it for such important meetings. If the low blame condition isn’t low blame to someone, then it is hard to know how to interpret their views on “low blame” situations.

8. Experiment 2 Participant #153

a. Quantitative Responses

- **Ought:** 50
- **Can:** -50
- **Blame:** -40

b. **Qualitative Response:** “Adams promised to meet Brown at noon - so he knew at 11:45am that he ‘ought to’ meet Brown, but because of assumingly unanticipated car trouble he could not do so.”¹⁵⁶

c. **My interpretation:** In one sense, this is a judgment about OIC in that participant mentions Adams’ inability to meet and affirms something about his “ought” to do so. However, there is an important nuance here that makes me think this judgment isn’t in response to OIC since participant writes that “[Adams] knew at 11:45am [*sic*] that he ‘ought to’ meet Brown,”¹⁵⁷ which expresses something like Adams’ knowledge *that* he ought to carry out his obligation at the time. I don’t think this is dealing with “ought” implies “can” because it sounds as if participant is merely affirming that Adams knows that he has an obligation at noon. This also brings in a fact potentially more corrosive to the data than any other mentioned so far: at 11:45, *we the readers* know that Adams can’t make it on time, but it is unclear that Adams can know this even if he has reason to think he might be late. This might explain why some participants make mention of alternate methods for Adams to make it—e.g., participant 184 mentions that Adams might take a taxi.¹⁵⁸ These participants might be taking on Adams’

point of view, which is first-person present and therefore not epistemically fated: Adams might be cursing at his car breaking down only to step out into the road right when a taxi slows to his feet. It is really only from the third-person omniscient perspective that we have knowledge that Adams can't make it on time.

9. Experiment 2 Participant #168

a. Quantitative Responses

- **Ought:** 23
- **Can:** -32
- **Blame:** 0

b. **Qualitative Response:** “After Adams' car breaks down, the delay makes it almost impossible for him to make the train on time. The blame is indeterminable given the info; it may be more Adams' fault or perhaps more the car manufacturer's fault. Adams ought to be able to meet Brown by noon, but given the change in circumstances, he cannot. What should be isn't always what can be.”¹⁵⁹

c. **My interpretation:** Again, what appears like a quantitative rejection of OIC—Adams is 23 “ought” and -32 “can”—is complicated by the qualitative response in which we learn that “the delay makes it *almost* [my emphasis] impossible for him to make the train on time”¹⁶⁰ and that “Adams ought to *be able to* [my emphasis] meet Brown by noon.”¹⁶¹ The ideal story breaks down somewhere, because these qualitative data don't reflect a judgment that squarely reacts to the details that press on OIC: it really must

be that participants see a fully robust impossibility facing Adams and nonetheless respond that he “ought” do precisely what that impossibility denies as possible action.

I get that Chituc et al. are trying to tease out the relationships between concepts of “blame” and “ought” and “can” and that they aren’t, at every turn, saying that these quantitative data are either supportive or negative to OIC, but these three concepts only bear on the OIC debate if we think that participants are really grappling with situations in which “ought” is pit against “can” *and* participants really are reacting to the “ought” how we intend they do and the “can” how we intend they do. If many of these participants are endorsing “ought” so as to express any of the following, then we just aren’t getting at anything identical to OIC: Adams is...

- accountable,
- blameworthy,
- expected in some general sense,
- expected by Brown,
- knowledgeable that he has an obligation,
- someone who should keep his promises,
- obligated to do everything in his power even if that means he fails in the end, and Adams is currently in a position where he has good reason to think that he will ultimately fail,
- etc.

In other words, teasing out how “blame” influences judgments of “ought” and “can” depends on us being reasonably confident that we are really getting “ought” and “can” judgments in just the way we care about them insofar as they relate to the Kantian principle. A host of interesting claims can emerge from the OIC study, those about how people discuss blame and use “ought” to convey

various judgments related to accountability and expectation, but maybe not many related to OIC. Or, simply put, if Chituc et al. want to discuss their results in light of OIC, they need epistemic justification in thinking that their data speak to OIC as the philosophers they challenge are thinking about it. Of course, there is one last experiment in the OIC study, OIC experiment 3, but I'll hold off on delving into it until chapter five when I discuss my own study since this is the experiment that I partially re-created. It is also the only one that lacks any qualitative data for this kind of analysis in the original OIC study.

Finally, we can wrap this chapter. In the previous chapter, on the Nahmias et al. work, we saw the MPEP in its many forms as it relates to experimental philosophy of free will and moral responsibility. This was done in part to show a subject-independence to the problem since free will taps into questions of metaphysics as much as it does ethics: pick your area in philosophy and I'd be willing to bet that the experimental work there is affected by the MPEP. And so, to preempt my own study, which will be explored in chapter five and beyond, we dove in and looked at the OIC study at a more fine-grained level, only to affirm the subject-independence of the MPEP and to see a more complicated picture of the MPEP. In fact, the more richly (i.e., qualitatively) we explore people's judgments, the more evidence we see of the MPEP existing and being a significant problem. I used language of the ideal story to highlight how many precarious moments in the collection of quantitative data exist and how they seem to often break down. It is often difficult to pinpoint precisely where the story breaks down, but it seems quite clear that *somewhere in here* or *likely over there* it did break down. When exploring my study in chapter five, we'll see that the picture is even more complicated than we've seen so far. And, given how complicated people's moral thinking is, why wouldn't it be? There we'll see how this complexity gives rise to serious doubts in x-phi but also to new ways forward.

OK, so now I'm moving on. But before getting to my study, one last task remains: justifying why I created my study the way I did and justifying its membership in the x-phi canon. The next chapter begins right there.

CHAPTER FOUR

Why I Crafted My “Ought Implies Can” Study the Way

I Did: All Roads Lead to Triangulation of Methods

I. The Situation in X-Phi Isn't Great, But a Move Toward Triangulation Will Likely Help

Admittedly, the picture I've painted of x-phi research in its current state looks bleak. The methodological problem of experimental philosophy runs deep and wide: for any given philosophical question, experimental philosophers might not be successfully eliciting or capturing people's genuine philosophical judgments *and* even if we could be sure said judgments were elicited and captured for any given philosophical question, we wouldn't be any more sure that we've succeeded in eliciting or capturing relevant judgments for other philosophical questions. One especially challenging feature of the MPEP is the fact that philosophical questions, even those about common concepts such as “knowledge” and “obligation,” can be deceptively complex and difficult to parse out. Consider, for example, the fake barn case that I have been referencing thus far in which someone sees what is actually a barn, notes that it is barn, but isn't aware that the barn stands near to a number of fake barns which are common to a particular stretch of land wherein the real barn sits, all of which leaves an outside evaluator to contemplate whether this someone truly *knows* that what they see is a barn.¹⁶² This scenario, a base case in epistemology, is just *not* straightforward. It is superficially *about* whether someone knows whether a structure is a barn or not, a simple and straightforward topic, and yet similar thought experiments in epistemology, such as Gettier's original cases¹⁶³ and the subsequent derivatives, rattled and continue to rattle philosophers due in part to the fact that many people's intuitions about knowledge clashed with accepted epistemological theory.¹⁶⁴ What's in a barn façade? A great deal, apparently.

In addition, some debates in x-phi don't seem like they can be resolved by the administration of more surveys. Some of the stalemates are about *which* surveys are or aren't effective in eliciting this or that judgment, and so no further surveys can settle the differences: if two philosophers don't agree on the rules of the game, then they surely won't agree on what data counts as a scored goal.

By looking at the experimental philosophy of free will in particular, we saw this bleak picture very clearly. In relation to this area of research, the MPEP is a problem that can't be ignored; there were simply too many plausible theories regarding Nahmias et al.'s data that haven't been accounted for—e.g., the theory that participants ascribe free will in almost all cases—such that we lack epistemic justification in concluding that the results from Nahmias et al. come straightforwardly from genuine compatibilist judgments participants have. In the x-phi of free will and moral responsibility, we also saw that a general over-reliance on quantitative surveys leaves philosophers in a number of these stalemates. If one group of philosophers understands the free will debate wherein compatibilism is a coherent position, and another group views compatibilism as incoherent, then we can't just throw more surveys at the problem in hopes of resolving deep-rooted disagreements. The message that came out of this analysis was that quantitative surveys alone can't deal with the MPEP, nor can they deal with the philosophical disagreements bound up in how x-phi surveys should be crafted. At times, I hinted at the solution being in employing qualitative methodologies, something I will pick up on in more depth in this chapter.

From there, we looked at the OIC study to see precisely where during a study the quantitative survey is prone to lead to unhelpful data, where the ideal story of data collection falls apart. By looking at brief qualitative responses by study participants in the OIC study, we saw that a good portion of them indicated that subjects weren't forming judgments based on the

philosophically relevant content in the thought experiments they read. This is the MPEP in its most clear form: based on these qualitative responses, which evidenced that subjects didn't always react to the relevant features of the thought experiments, we should accept that we have no good justification for thinking that the quantitative data from the OIC study reflect the kind of judgments they are supposed to reflect. This fine-grained look at the OIC study revealed not only where the ideal story might be breaking down, but also where we can call on qualitative methodologies to better understand *what* transpires when a subject completes a classic, survey-based x-phi study (which is what I pick up with my own research in the next chapters).

Despite this problem-ridden landscape in x-phi research, I'll suggest a possible solution in this chapter and realize it in subsequent chapters: the use of qualitative methodologies and the triangulation of methods. My goal in this chapter is to point to the viability of certain methodological practices as a solution to an over-reliance on quantitative surveys, thereby setting the stage for my own study which draws on certain qualitative methods. In the next chapter, I'll unpack the details of my own study and start getting to my findings.

On the assumption that readers know what qualitative methods are in general, I'll move to the term that I was initially unfamiliar with when conducting my literature review: triangulation of methods. So what is triangulation in the context of research methods? In short, it is the practice in qualitative research of using multiple experimental methods to home in on a particular object of study or phenomenon.^{165, 166} It gets its name from a practice far removed from social science research; specifically, "the concept of triangulation is borrowed from navigational and land surveying techniques that determine a single point in space with the convergence of measurements taken from two other distinct points."¹⁶⁷ And while there might be a perfect, or even acceptable, way to analogize from its use in navigation to research methods, wherein one research

methodology functions like one essential measurement for triangulation and another methodology functions as another, I will abstain from such precise analogizing because the term “triangulation” is employed more loosely in the research methods context: “The basic idea underpinning the concept of triangulation is that the phenomena under study can be understood best when approached with a variety or a combination of research methods.”¹⁶⁸ Even when I move to a more specific discussion of what is called “triangulation of methods,” which is a distinct kind of triangulation that is achieved by using multiple research methods rather than multiple theoretical frameworks or multiple investigators, I am using the term as it is understood in the context of research methods, which doesn’t rely on a strong analogy between navigational triangulation and constructing a study in the social sciences.¹⁶⁹ As I’ve just iterated, the idea of triangulation in the research methods context is coherent even when used in a general sense to convey the study of something by making use of multiple methodologies. It is in this general sense that I want to bring the concept of triangulation into experimental philosophy, because, as far as I know, it hasn’t been robustly developed or fully articulated in that context, and I think there are good reasons for thinking that introducing triangulation to x-phi would be helpful for learning about how people think about philosophical issues.

Before I look at how triangulation of methods might operate in the x-phi context, I do want to elaborate on precisely what assumption I am making in doing so. Ultimately, I am invoking triangulation in a rough and unrefined sense that borrows its central idea from its use in research methods. Rather than this lack of refinement being a problem that halts my project, I take it to be somewhat of a necessity given that my project here is largely exploratory: I’m trying out some kind of triangulation of methods by actually doing a study that involves it and analyzing my results across multiple methodologies on the assumption that they can each speak to the unique ways that

people think about philosophical issues in ethics and beyond. So yes, the very notion that multiple methodologies could tell us something about the same phenomenon, in this case a layperson's philosophical judgment, is an assumption I make in this project (and it strikes me that it is the core assumption for any application of triangulation of methods). I acknowledge that there is some risk in this assumption given that there is debate over how different research methodologies can speak to each other given their unique epistemic assumptions and constraints.¹⁷⁰ However, given the nature of philosophical judgments and the methodologies I applied, I thought the assumption was generally safe to make, especially as an exploratory project that might prove unsuccessful as easily as it could prove successful. Specifically, I am interested in people's philosophical judgments which are, straightforwardly, how they respond to or think in regard to certain philosophical issues; they are not necessarily subconscious or non-conscious cognitive or neural events that people are fundamentally unaware of. Therefore, when I ask people to articulate their judgments by speaking aloud their thoughts while completing a survey and when I subsequently ask them to answer targeted questions about their thinking, I am employing methodologies that straightforwardly get at their philosophical judgments. There are certainly limitations to these methods in isolation and in combination, but they aren't so grave as to prevent exploratory research that makes use of them, just as I've done. In other words, it seems generally plausible that the think aloud method, the interview, and the quantitative survey each can give us insight into how people think about philosophical issues, despite the limitations of each methodology and regardless of whether we have—yet—a precise account of what epistemological framework undergirds the integration of data from a survey with data from a think aloud session, etc. (Also, as I explore in detail in the next section of this chapter, the work and words of other researchers point to the viability of incorporating qualitative methods like the interview and think aloud method into the x-phi

context.) And so I tried out a triangulation of methods in the x-phi context and I argue, throughout this dissertation, that it was successful; I will try to convince the reader of this as well. From here, I hope that future research will pick up on refining precisely how triangulation of methods operates in the x-phi context, and how researchers can more fruitfully come to understand philosophical judgments and the methods that can accurately capture them.

As I indicated, I've yet to encounter a robust use of triangulation in experimental philosophy that takes full advantage of qualitative research methodologies. Experimental philosophers typically craft surveys based on what they and other philosophers think will be valid (or they draw on famous thought experiments in the philosophy literature), administer their surveys, then create new surveys if *they* think the original data or something else reveals some kind of potential flaw in the original survey. That is, experimental philosophers aren't doing much in terms of diversifying their methodologies, nor are they doing much in terms of validating their quantitative surveys. This should strike philosophers as particularly odd, given that the way philosophers use thought experiments such as those used in x-phi research is almost never in a wholly quantitative fashion. Before a lecture hall, or in conversation with others, the philosopher—in my experience—never presents a famous thought experiment and then requests that their interlocutors reply only in a fashion that could be captured by a quantitative measure. For example, the philosopher doesn't ask of students, 'On a scale of one to seven, note your agreement to the following statement, and *please* don't elaborate or explain your response: it is morally permissible to divert the trolley.' In almost all cases, the thought experiment is introduced *so as* to promote a full, robust dialogue or conversation on the issues at hand. One interlocutor provides a response to the thought experiment, and perhaps their reasoning for their selection, to which the other responds by questioning features of the first's response and perhaps by presenting

a slightly modified version of the same thought experiment to explore the specific contours of the first's initial response and reasoning. That is, the employment of thought experiments typically has a *qualitative* feature built right into it: it creates an ever-evolving dialectic that otherwise might not have begun.

And this is true even when one philosopher introspectively reflects on a thought experiment. She doesn't report a quantitative response to a thought experiment without considering her reasons, reflecting on possible counters to her thinking, etc. In a sense, she follows through with an internal, reflective dialogue of one. And yet, the preferred methodology for experimental philosophy has been—for reasons I might pin to convenience, imitation of contemporary psychological research which has some kind of scientific clout, and a suspicion for data that feel 'subjective' in nature—the *quantitative* survey, a tool which ultimately fails to capture the essence of the employment of philosophical thought experiments in the real world. Unlike many psychologists, experimental philosophers seem satisfied relying solely on a singular approach to understanding how people think. And yet we know how nuanced people's views of ethics and beyond can be.

The problem I am highlighting isn't that experimental philosophers use quantitative surveys to gather data, but rather that they use them (nearly) exclusively. To avoid this kind of problem in their respective domains of inquiry, many social scientists have made strong efforts to use multiple methodological approaches in coming to understand an object of study. Just as a detective looks to physical evidence, eye witness accounts, subjective reports regarding aural data, etc., so too does many a researcher attempt to home in on certain events—psychological or sociological ones, for example—by approaching them from multiple angles, using multiple methodologies. So, rather than just advocating for using qualitative methodologies as such, I

advocate using them as a means for achieving a triangulation of methods. In short I endorse what's at the heart of triangulation: "When designing and conducting research, qualitative investigators frequently combine methods such as interviewing, surveys, and observation across variable times and in different places in order to collect data about their research phenomena from multiple perspectives and in different contexts."¹⁷¹ If x-philosophers looked at free will/moral responsibility or "ought implies can" through multiple methods, then it seems clear that they would learn more than if they relied solely on quantitative surveys which have just as many limitations as they do advantages. This is not to say that qualitative methods don't also have limitations, but their use in combination with quantitative surveys can shed more light on a phenomenon than either methodological approach could in isolation: just as I don't advocate for relying solely on quantitative methodologies, I am *certainly not advocating* for x-philosophers to use qualitative methodologies exclusively.

Outside of x-phi, use of multiple data-gathering methods is commonplace, and for good reason. For example, in a mixed-methods study on substance use in people with psychotic disorders, researchers had some subjects complete semi-structured phone interviews in addition to quantitative questionnaires.¹⁷² Here the researchers summarize the benefits of their approach:

To allow for the development of appropriately targeted interventions for co-occurring substance use and psychotic disorders, it is important to understand why people with psychotic disorders use substances and where reasons for use differ between substances. This study was uniquely placed to answer these questions, because it is only the second study to directly compare reasons for tobacco, alcohol, and cannabis use and the first to supplement quantitative data on patterns of reasons for substance use with rich qualitative descriptions of individuals' reasons for use. This approach allowed us to gain a greater

understanding of reasons for substance use among people with psychotic disorders and led us to suggest that people with psychotic disorders use tobacco, alcohol, and cannabis for different reasons.¹⁷³

In this case, the qualitative interviews allowed subjects to further articulate their reasons for substance use. Were researchers to have only used quantitative surveys, it is probable that they would have missed out on important information about the focal point of their research. This is the *kind* of triangulation that I want to advocate for in x-phi research because it seems especially appropriate when trying to get at philosophical judgments/intuitions, which strike me as particularly opaque. Also, this mix of methods would allow for philosophers to invite participants to explain the reasoning behind the judgments they provide in relation to a thought experiment, which is clearly helpful in understanding those judgments. Of course, this invitation to provide reasons can, in the course of an experiment, occur after participants give initial quantitative responses so that the request for explanation doesn't interfere with their initial thoughts about a particular thought experiment or follow-up question. Ideally, we want to know not only *what* judgment a person forms, but also *why* they came to that judgment; triangulation of methods within the context of x-phi can help with this project.

Not only can qualitative methodologies produce data that informs or adds to quantitative data, it can also provide information that reveals inaccuracies or incompleteness with something studied solely or largely by quantitative means. For example, sociologist Phil Zuckerman thought there might be more complexity to how we understand people's religious (or non-religious) identities after studying statistical data from quantitative surveys. In particular, he noticed interesting tensions in data on Scandinavians, who, in surveys from 1990-1993 on beliefs and values, were reported to express the *lowest* belief in sin¹⁷⁴—a core tenet for many Christians—

while other data listed Scandinavians as being predominantly “Christian.” Compared to 89% of Americans who reported belief in sin, an unsurprising figure given the high population of Christianity in the U.S., just 24% of Danes said they believe in sin.¹⁷⁵ And yet, global data from that same time, 1993, has Denmark as 92.8% Christian,¹⁷⁶ and data from the CIA’s “The World Factbook” currently lists 74.8% of Danes as Evangelical Lutheran.¹⁷⁷ This prompted Zuckerman to learn more about Scandinavian religious identity, which he discovered to be more interesting and complex, by conducting a series of in-depth interviews with Scandinavians about their religious beliefs.

Consider two brief examples Zuckerman introduces at the start of his book *Society Without God*, which details his exploration of religion and non-theism in Scandinavia. First is a man from Denmark who identifies as a Christian while disbelieving in Jesus’ divinity, heaven and hell, and the revelatory status of the Bible—as for God’s existence, he’s not sure what he believes.¹⁷⁸ He explains “[he’s] *believing in good things in human beings, which are the real things of Christianity,*” giving examples of helping the elderly and abstaining from killing.¹⁷⁹ Second is a woman from Sweden who identifies as Christian yet disbelieves in the bread and butter of Christian theology: God, Jesus, Bible, heaven and hell.¹⁸⁰ In her words, being a Christian means “*to be a decent human being and respect other people and, yeah, to be a good person.*”¹⁸¹ In total, Zuckerman learned that many Scandinavians have a notion of Christianity very different than we might expect were we to simply review broad survey data. In fact, he notes that what these two interviewees express is squarely in line with secular humanism, not Christianity as typically conceived in America.¹⁸²

Of course, we can imagine crafting a number of quantitative questions that would have similarly revealed what Scandinavians took to be Christianity—e.g., ‘As a Christian, do you

believe in God?’—but that hardly negates the value of Zuckerman’s use of qualitative interviews. It wouldn’t have been clear what specific questions a revised quantitative survey would need to ask of these Scandinavians without Zuckerman’s interviews in the first place. That is, he unearthed what kinds of questions would need to be asked of the population in order to reveal more nuanced information about their atypical notion of Christianity. Also, even with a revised quantitative questionnaire, it is difficult to imagine that all of the nuance that Zuckerman was able to record in interviews would be well-captured. For example, the two interviewees above connect being Christian with being a good person, which places moral behavior as central, whereas other Scandinavians connect a Christian identity with other factors, such as association to the church and to some larger cultural identity. One elementary school teacher articulates that she views herself as a Christian which “*means that [she] respect[s] the church and if [she] ever get[s] married [she] want[s] to definitely get married in the church*”¹⁸³; one preschool teacher who doesn’t believe in God but views himself as a Christian explains that “*it’s because people in Denmark are Christians. And it’s a term you use when you speak with people from other countries. So people often ask you—are you religious or what’s your religion?—then it’s easy to say I’m Christian.*”¹⁸⁴ It seems clear that the format of a qualitative interview allows us to get an accurate sense of this kind of thinking regarding religious identity in a way that a quantitative survey can’t. And yet, Zuckerman’s interviews alone don’t provide the whole picture: he was able to interview a handful of Scandinavians, but quantitative surveys can more easily gather data on many more people. In addition, were Zuckerman to lack the original quantitative data, he might not have realized that there is something more interesting going on in Scandinavian religious identity than meets the eye. Through a kind of triangulation of methods, though, Zuckerman was able to add to the accuracy and complexity of the picture detailing the religious identities of Scandinavians: both the

qualitative and the quantitative data can inform each other and the researcher regarding the way certain people engage with religious ideas and identities.

These examples of triangulation of methods (of a kind), the first from a study on substance abuse and then one on religiosity, describe research areas outside of x-phi. It is natural, then, to wonder what such a methodological practice might look like *in* x-phi. Unfortunately, there just isn't much to draw on in the x-phi literature. In conducting a literature review, I found (at the time of investigation) only one x-phi study which had taken advantage of a kind of triangulation of methods: one by Joshua May in which he investigated what emotional reactions people tend to have toward human reproductive cloning.^{185, 186} In May's study, participants read a brief background text about human reproductive cloning before they were asked to respond to a mix of quantitative and qualitative questions designed to get their moral and emotional reactions to human reproductive cloning. The first of two qualitative questions during the study asked participants to "list any emotional reactions" they had to reading about this cloning "whether positive or negative"¹⁸⁷; the question was strategically placed prior to the next quantitative questions, which asked participants to select emotional reactions from predetermined lists, which "provided an opportunity for subjects to report what they felt immediately after reading about cloning and without being biased by a preset list of choices about what to feel."¹⁸⁸ With the second qualitative question, which concluded the study, "participants were allowed to provide any feedback on the survey. They were specifically encouraged to identify any emotions they felt towards human reproductive cloning that were not on the lists."¹⁸⁹ May's study is focused largely on the emotional reactions participants had, though he also asked two brief questions as to whether they think human reproductive cloning is immoral or should be illegal¹⁹⁰; this is somewhat different from a hunt for something like the philosophical judgments/intuitions that most experimental philosophers are

after. May doesn't, for example, pump subjects' intuitions about whether cloned humans participate in personhood by presenting subjects with cleverly invented cases. Still, the study fits squarely into the x-phi project since May uses his findings to discuss how people think about cloning and whether there is widespread intuitive support for such cloning being disgusting, a key premise in one form of the "argument from repugnance," in which one appeals to disgust as demonstrative of cloning's immorality.¹⁹¹ What is valuable in this study is that May's use of mixed-methods provides a kind of triangulation of methods to better capture participants' emotional reactions to the background text on human reproductive cloning. While this is a very small-scale example of what triangulation might look like in x-phi—i.e., May doesn't conduct full interviews or something similar—it seems clear that the mixed-methods approach provides more information than could have been recorded were subjects to respond solely to quantitative questions.

So where does this leave experimental philosophers? It leaves us in a particularly difficult, but manageable situation. We need to be confident that x-phi is getting reasonably close to capturing people's relevant philosophical judgments or intuitions, and triangulation seems to be a good way forward in trying to do so. A vignette can't just *seem* to tap the right judgments/intuitions from a philosopher's or a psychologist's point of view; it has to elicit the right kind of judgment/intuition *and* researchers need to have evidence that such a judgment/intuition has been captured. To be able to claim that people from culture *x* differ in their philosophical judgments regarding incompatibilism when compared to people native to culture *y*, it *must* be the case that data from each culture reflect genuine philosophical judgments regarding incompatibilism. It is not enough to say, 'When given the supercomputer vignette, people from culture *x* and *y* produced significantly different responses.' Although, this could be the starting

point for further study of how genuine these philosophical judgments are. But it isn't difficult to imagine that people from culture *x* or *y* or both tended to produce responses that didn't truly track genuine philosophical judgments—and this hypothetical is plausible for at least two related reasons.

First, this hypothetical is plausible when we consider both the difficulty philosophers find in crafting vignettes that meaningfully elicit judgments of incompatibilism (or other philosophical ideas) *and* the fact that certain ideas of freedom, choice, and moral responsibility might turn out to be culturally specific. Second, this hypothetical is plausible given the complexity of certain philosophical questions. While the complexity of the doctrine of incompatibilism may seem unique in making the job of the experimental philosopher difficult, I think it's probable that most if not all of the questions that x-phi researchers are likely to puzzle over are sufficiently complex or perplexing such that judgments in response to them aren't so easy to elicit and identify. This is because the things philosophers debate and write often about tend to be things that are on their face complicated and controversial—e.g., questions of free will/moral responsibility or the mind-body problem—and if they happen to be simple on their face then the debate is about some relatively perplexing feature of them—e.g., questions as to why it is so difficult to define knowledge (a seemingly straightforward thing).

Fortunately, everything I've been discussing in this chapter as a challenge for x-phi leaves experimental philosophers with a potential solution in qualitative methodologies and triangulation of methods. It's not an easy way forward, but it is one. We are not bound by surveys alone, and we can take creative measures to learn more about how people think about philosophical questions and why they choose the quantitative responses that they choose. For any given philosophical question, if every use of qualitative methods reveals that people really are providing genuine

judgments or intuitions then we have epistemic justification in making claims about people's quantitative data as more or less reflective of how people think about that question. But as we've seen, based on just a preliminary examination of the OIC study, this is unlikely to be the case. It seems far more likely that philosophical judgments are difficult to get. But engaging with a given philosophical question using triangulation enables us to point to *what* we feel confident claiming we've successfully captured in people's thinking, and *what* we think we still need more insight into. The solution, then, isn't to use qualitative methodologies *instead* of quantitative ones, but to use both in chorus to approach the projects of x-phi from multiple angles—i.e., triangulation of methods. We have every reason to be skeptical that the current x-phi program overly relies on quantitative surveys, and in the remainder of this chapter I will argue that we also have good reason to think that qualitative methods will be beneficial. Then, in the next chapters, I will outline my own study to show what came of my attempt to incorporate qualitative methodologies into the x-phi research program.

II. There is a Lack and Prospect for Using Qualitative Methods in X-Phi

In this section, I explore the relationship, or rather non-relationship, between experimental philosophy and the use of qualitative methodologies. I make the case that other academic disciplines similar to x-phi make good use of qualitative methods, specifically the ones I used in my study (the think aloud method and the qualitative interview), and that there is good reason to think that qualitative methods could be useful in certain x-phi studies. This chapter in total, including this section, is a primer for outlining the specific design and results of my own study. Before I get to it in full detail in the next chapter, I show here and throughout the rest of this chapter that I had good reason to craft my study as I did, and I also show that even if my study falls short

in many ways *there is still good reason* for experimental philosophers to incorporate qualitative methodologies into their future studies.

To begin, I do need to outline the basic thinking behind my study. I developed my study because I wanted to learn more about philosophical judgments by “observing”—as best as possible—the thoughts and considerations that lead experimental subjects to their decisions on quantitative surveys. To do so, I conducted an experiment based on OIC experiment 3 in which I had subjects self-report these thoughts in detail *as* they are reacted to thought experiments and *afterwards* in an open-ended interview session. My thinking was that this original experiment by Chituc et al., and by proxy experiments similar in kind, did not capture subjects’ coherent, potentially divergent judgments. Rather, the experiment puts the label “judgment” on what is more likely an inaccurate quantitative value arrived at by inorganically funneling subjects’ non-uniform cognitive responses through limited survey responses. In other words, my thinking was that the ideal story of x-phi data collection breaks down at some point in a significant way. If I ended up being more or less correct, then I would have shown that (many) x-phi studies are not capturing what they are said to be capturing. If, however, the ideal story is corroborated, then my research would strengthen the claims made by experimental philosophers and other researchers who employ similar methodologies.

To carry out my research, I used two qualitative methodologies, the think aloud method and the qualitative interview, because I believed they could plausibly uncover phenomena obscured by quantitative survey questions. Before discussing the specific features of these methods, which I do in the very next section of this chapter, I think it valuable to first discuss the relationship between qualitative methodologies and experimental philosophy in general. Most empirical research in experimental philosophy is currently gathered using quantitative methods

rather than qualitative ones, such as interviews or the think aloud method. James Andow—from whose work I got the idea to use the think aloud method in the first place¹⁹²—notes that “experimental philosophers do not tend to use qualitative methods,”¹⁹³ and he demonstrates this by referencing a list of articles from 2014 on The Experimental Philosophy Page in which one empirical article out of thirty-six was labeled as qualitative—a label only given to articles “if qualitative data was collected, analyzed, and featured in the presented results (as something other than a manipulation check).”¹⁹⁴ While this data is solely from 2014, it gives us some sense of how rare qualitative studies are in current x-phi research. And, when considering interviews and the think aloud method as particular qualitative methods, my research of the entire PhilPapers Experimental Philosophy database (at the time of completion) revealed that no study currently within this literature specifically uses either method to collect data for philosophical analysis.¹⁹⁵

Though these two qualitative methods are not found squarely within x-phi, they are employed in related fields, a fact which lends credence to their potential applicability within x-phi. For example, interviews of a kind are used in certain branches of philosophy, typically those that intersect with the study of ethics in one way or another—e.g., bioethics, business ethics, and military ethics. This is promising in that my project is one of ethics, too; ethics finds itself naturally, then, at the bridge between x-phi and other philosophical work on empirically informed issues. The authors of one article in bioethics entitled “Reexamining Healthcare Justice in the Light of Empirical Data” describe in the abstract of their work that they “*undertake an empirical analysis of the concepts of justice held by healthcare personnel, gleaned from a qualitative analysis of interviews on the subject of ethical dilemmas in everyday practice.*”¹⁹⁶ In a publication in business ethics, “Consumer Ethics in Japan: An Economic Reconstruction of Moral Agency of Japanese Firms – Qualitative Insights from Grocery/Retail Markets,” the researcher writes that “an interview

procedure”¹⁹⁷ was used, and the author of a publication in military ethics “Should Soldiers Think before They Shoot?” describes in the abstract that the publication’s “*empirical materials are qualitative and build mainly upon extracted information from interviews and informal conversations with Norwegian soldiers and officers serving in Afghanistan under ISAF’s Regional Command North in 2007 and 2008.*”¹⁹⁸ In addition to research in these areas of ethics, interviews have been used in other fields as well, including social philosophy.¹⁹⁹ While the goals of all these disciplines are not necessarily aligned precisely with those of x-phi, it is clear that interviews have been effectively used in philosophical investigations not too far removed from those in x-phi. Even these examples and my current project, in their similarities, blur the boundaries between sub-disciplinary goals.

Interviews are also used in psychology research similar to x-phi experimentation (again the boundaries of x-phi start to fade upon closer inspection). A well-known example is from Jonathan Haidt, a moral psychologist, who uses interviews of a kind in his research described in his book *The Righteous Mind: Why Good People Are Divided by Politics and Religion*.²⁰⁰ For example, in order to capture people’s thoughts about hypothetical scenarios in which no one is harmed yet a sacred item is disrespected, participants were asked about a woman who decides to clean her bathroom with an American flag.²⁰¹ In another well-known example, Carol Gilligan, in her work *In a Different Voice*, draws from interviews with women on topics of morality and moral conflict while developing her ideas on the different way that many women reason about morality.²⁰² The work of Haidt and Gilligan are so similar to work in x-phi, and my current project, in that they use empirical data to make claims about the ways people think about moral issues, that they might be considered part of its historical canon in some sense. I think their examples, based heavily in qualitative methods, support the case for more qualitative methodologies.

Though Haidt and Gilligan might count for interview use in the x-phi movement, there have been no experimental philosophy studies (to my knowledge) that make extensive use of the think aloud method.²⁰³ However, other empirical research related to philosophy has. For example, in one study, entitled “Judging the Plausibility of Arguments in Scientific Texts: A Student-Scientist Comparison,” the think aloud method was used to answer questions related to scientific understanding and epistemology.²⁰⁴ In describing their work, the researchers write,

Using a think-aloud procedure, we identified the specific strategies students and scientists apply when judging the plausibility of arguments and classifying common argumentation fallacies. Results indicate that students, compared to scientists, have difficulties forming these judgements and base them on intuition and opinion rather than the internal consistency of arguments.²⁰⁵

Similar to my research, this study seeks to capture how someone thinks about a given philosophical issue, in this case a potentially fallacious argument. In addition, the researchers include a retrospective interview component in order to “gain further insights in the procedural and declarative knowledge used by scientists and students during task completion” and they use interview probing methods to “assess clarity of instruction, perceived task difficulty, familiarity with argumentation fallacies, and self-reported strategy use.”²⁰⁶ As with my own study, this one employs not only the think aloud method to gather data on people’s thinking, but it also uses the interview as a means of clarifying people’s thinking and other important information.

In another study, entitled “On Imagining What is True (and What is False),” researchers made use of the think aloud method in order to capture how subjects reasoned through logic tasks that centered on assertions which included the connectives “and,” “or,” and “if.”²⁰⁷ In their discussion of the think aloud data, the researchers, in reference to how subjects reasoned during

the experiment, report that “what the participants said as they thought aloud was revealing.”²⁰⁸ While questions of logic are unique within the larger discipline of philosophy in that they more closely resemble questions of mathematics rather than those of ethics or metaphysics, this example shows that my use of the think aloud method is not far removed from another philosophical investigation that found the method to be fruitful.

In addition to the effective use of interviews and the think aloud method in other subfields of philosophy and those adjacent to it, some philosophers have explicitly claimed that there could be value in these two methods specifically, and qualitative methods more generally, for x-phi research. For example, Womack and Mulvaney-Day argue for the use of interviews in experimental philosophy research given the limitations they see with the survey method used by experimental philosophers.²⁰⁹ Specifically, with reference to the surveys used in x-phi, they first claim that “given the abstractness and brevity of the scenarios used, survey response patterns underdetermine the psychological processes leading to the responses.”²¹⁰ In addition, they claim that “employing methods that do not allow respondents to reframe scenarios to reflect features that they find salient in fashioning their responses may cause philosophers to miss out on potentially fruitful explanations for responses.”²¹¹ They therefore recommend that “philosophers consider including qualitative methods and expanded interviewing processes. These in-depth approaches can elicit the meaning and the decision-making processes underlying the respondents’ answers to the questions.”²¹² On their view, experimental philosophers could use qualitative methods as unique tools that can reveal information lost in quantitative surveys.

Philosophers Elizabeth O’Neill and Edouard Machery also note the potential value of using interviews:

Expanding the kind of questions addressed within experimental philosophy requires expanding the range of experimental methods to which experimental philosophers appeal. They should avail themselves of the complete range of experimental methods used in the sciences. Nonexperimental empirical methods such as interviews and qualitative studies could also be fruitfully used by experimental philosophers.²¹³

And so do Nahmias et al. in their publication “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility” when discussing the limitations of their work on free will, such as the supercomputer thought experiment. Specifically, they note “the method of surveying judgments in response to scenarios...should be combined with other methods, including coding participants’ responses in more open-ended interviews about the relevant concepts,” which aligns nicely with the study I conducted on “ought implies can.”²¹⁴ In another publication on free will, Nahmias et al. emphasize the value of the “‘talk aloud’ method”²¹⁵—which they describe in terms identical or nearly so to what I understand as the think aloud method—and of a type of interview, namely the “phenomenological interview,” which is characterized by “extensive interviews with subjects about particular experiences, treating the subjects as co-investigators of the phenomena.”²¹⁶ After outlining a lack of effective research on the phenomenology of free will, they offer support for such practices in presenting their own research on free will which uses each of the two techniques. Specifically, they conduct a talk-aloud study which has participants talk through their deliberations and choices in selecting a hypothetical apartment, though Nahmias et al. make very little of the findings and present almost nothing of the transcripts.²¹⁷ And they also conduct some phenomenological interviews, but they report having not coded them and they present no material from the interviews.²¹⁸ Nonetheless, the exercise of these methods serves as some kind of support for their value in that the authors treat their use as the preliminary part of a project that more

researchers can get involved in rather than methodological dead ends.²¹⁹ (On first read, this study sounds quite similar to the one I conducted in that subjects are asked to talk aloud and their utterances will be taken to bear on a philosophical issue; and there is also an interview involved. At one point in this publication, Nahmias et al. even highlight the value of a kind of triangulation, which I will discuss at length later, that would involve connecting reports from laypeople along with “behavioural and neuropsychological data.”²²⁰ However, the study turns out to be more about *what it is like* to deliberate and make a choice—i.e., an investigation into the phenomenology of free will—rather than a study designed to get at *what people think is true* regarding theories of free will. So, while I think it important to view it as work in x-phi, the brief treatment of the qualitative results and the goals of the study render it more or less incomplete as an x-phi study that makes full use of the think aloud method, interview, or triangulation.)

Philosopher James Andow argues for the potential value of qualitative methods as he notes the unique benefits they provide:

What are the advantages of employing more qualitative tools? Qualitative methods can be used to gain insight into how participants think and talk about philosophically interesting phenomena for themselves. Of course, quantitative surveys can do this. However, qualitative methods can give researchers much deeper insights. The reason is that they remove certain barriers. For example, it might be otherwise very unclear how participants understand the question they are asked. Or there might be very important aspects of how a given participant thinks about an issue which a fixed response format might provide no ability to express, and which will consequently remain unknown to philosophy. Another barrier which is removed is that participants are able to respond in their own words, so one can gain an insight into the way in which the “philosophical” usage of a term, as employed

in a philosopher's experimental materials, may not match the ways ordinary folks use the term.²²¹

Andow later mentions interviews as one such qualitative tool.²²² And, importantly, Andow specifically advocates for the use of the think aloud method.²²³ In fact, it was this suggestion by Andow that prompted me to look into the think aloud method as a research tool. I also took his suggestion to heart when designing my study since I agreed with his optimism toward qualitative methodologies.

In sum, experimental philosophers are not currently using interviews or the think aloud method in their research (unless you count certain endeavors just outside the very specific label 'x-phi' or those that make minimal use of such methods), but researchers in other areas of philosophy are.²²⁴ With reference to x-phi, some philosophers advocate for the use of qualitative methods, or at least acknowledge the potential value in using them. This picture presents experimental philosophers with the opportunity to explore the value of interviews, the think aloud method, and other qualitative methodologies. And, that reality is what prompted me to craft my study as I did. Importantly, though, even if my study fails—something the reader can be the judge of in the next chapters—I think there is still good reason to pursue qualitative methodologies for just the same reasons as I did plus the additional reason that either 1) I successfully showed that it is valuable in my own study thereby incentivizing others to conduct similar studies, or that 2) I successfully showed *how not* to do it thereby incentivizing others to improve upon my failings.

III. The Unique Benefits of the “Think Aloud Method” and Qualitative Interviews

Having provided a general discussion of qualitative methods—and a brief vote of confidence for the think aloud method and the interview in particular—I will move to a very

specific justification of why I think the two proposed methodologies I used in my own study are especially helpful for x-phi research. Specifically, I argue they are good for creating a triangulation of methods with respect to studying people's philosophical judgments.

While most are familiar with the structure of an interview, many are likely unfamiliar with the think aloud method (just as I was prior to developing my study). According to Maarten W. van Someren, Yvonne F. Barnard, and Jacobijn A. C. Sandberg, whose book *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* is foundational for my present research, “the think aloud method consists of asking people to think aloud while solving a problem and analysing the resulting verbal protocols,”²²⁵ where spoken and written protocols are, respectively, “the speech and writings” that result from asking a person to think aloud during a task.²²⁶ Though the method is often applied to problem-solving in a very literal sense—e.g., problems related to architectural design²²⁷ and arithmetic²²⁸—I believed it could be applied to “problems” of philosophy presented in the form of thought experiments.²²⁹

Before I detail my use of the think aloud method for philosophical problems, I want to trace the method's history and applications. According to van Someren et al., the think aloud method emerges from the psychological method of introspection.²³⁰ Understood in conjunction with retrospection—in which “the subject solves a problem and is questioned afterwards about the thought processes during the solving of a problem”²³¹—introspection involves planned interruptions made by the subject for the purpose of reporting and interpreting his or her thinking.²³² According to van Someren et al., “introspection is based on the idea that one can observe events that take place in consciousness, more or less as one can observe events in the outside world.”²³³ On this view, the data of introspection are the conscious events of the person introspecting, which van Someren et al. describe as leading to both theoretical problems and

methodological ones. A core theoretical problem hinges on the question of how exactly introspection, itself a conscious process, is understood in relation to the (apparently distinct) events in consciousness that are being examined through introspection—for example, it's unclear if one can apply introspection to his or her own introspection.²³⁴ Another major problem is methodological in nature. Because introspection treats the events of a person's consciousness as the empirical objects of observation, there is no way for anyone but the person introspecting to access them; the unhappy result is a limitation on what scientists can appeal to and analyze as data for understanding how people think about certain things or tasks.²³⁵

The impact of such problems has been a wariness toward things similar to introspection.²³⁶ And the think aloud method certainly appears similar to introspection on the surface. Both involve some kind of systematic examination of a person's thought processes by way of asking that very person to report information, in one way or another, about what is occurring in his or her conscious experience during a given period of time or during a given task. Suspicion toward introspection, then, seems to justify suspicion of the think aloud method, because researchers are relying on individuals to provide accurate reports of their conscious experiences. However, the think aloud method bypasses such suspicion for two reasons: 1) the subject talks out what he or she is thinking *without* adding interpretation; and 2) the resulting verbal protocols become objective data that anyone can access, not just the subject.²³⁷ In short, verbalization does not interfere, or at least not significantly so, with the cognitive processes that introspection necessarily interferes with. In addition, the recorded protocols, now accessible to any researcher, can even help researchers detect if a participant is experiencing substantial interference from thinking aloud because they capture if a participant is struggling to think aloud or expressing that they are having trouble doing so. The think aloud method certainly has flaws, but none as problematic as those plaguing introspection.

Having unpacked the history of the think aloud method, I want to justify my use of it in the context of x-phi which differs in significant ways from the problem-solving contexts described by van Someren et al. The think aloud method is used to “observe” someone’s thought processes as they solve a problem; in line with this, I employed the method to “observe” subjects’ thought processes as they solved philosophical problems. But of course, philosophical problems are not like other problems: outside of philosophical logic, they rarely have definitive, universally (or near-universally) agreed upon solutions. This appears to be a problem for my present research since think aloud protocols are sometimes analyzed based on how a person arrives at a *correct* answer. For example, van Someren et al. describe an arithmetic problem for which there is one correct answer: “*A father, a mother and their son are 80 years old together. The father is twice as old as the son. The mother has the same age as the father. How old is the son?*”²³⁸ By having people think aloud when solving problems like this, researchers can recognize different strategies people take in coming to the correct response, 16.

With other problems, subjects are required to come up with a response that correctly satisfies a set of criteria. For example, van Someren et al. describe a problem given to architects in a study conducted by Ronald Hamel in which architects were required to design a facility that met certain criteria; Hamel wanted to learn how architects reasoned through the design process.²³⁹ Though criteria-based problems can potentially have multiple correct solutions, researchers can say whether or not they are successfully solved based on whether a subject’s response met certain criteria. However, when dealing with thought experiments in x-phi, there are often no universally agreed upon “correct” answers. This is true in two important senses. First, many thought experiments do not have a universally agreed up on correct answer in any mathematical or criteria-based sense because many philosophical issues deal with value judgments (i.e., it might be true

that having the train go down the track occupied by one person in the trolley problem²⁴⁰ preserves greater net happiness in that five lives are spared at the death of one, but one cannot hold that valuing net happiness *is* the correct way to perceive ethical decisions). Second, thought experiments used in x-phi are often designed such that *any* answer one provides satisfies the question being asked, assuming that the answer is one of the available choices on a survey. For example, reporting either that it *is* or *is not* morally permissible to have the trolley kill one person rather than five is a satisfactory answer since the thought experiment, at least in the context of x-phi, is designed to capture *what* one thinks is the best response, not whether one can “solve” a problem. On the face of it, then, the think aloud method appears to have little use in x-phi.

However, closer examination reveals that the think aloud method should serve x-phi well. According to van Someren et al., “the purpose of the [think aloud method] is to obtain data about a cognitive process.”²⁴¹ And therefore, despite the significant differences between typical use of the method and my research, the same core principle and application is preserved: the method gives researchers insight into how someone is *thinking through* some problem. And while I am not interested in some problem being “solved” or some criteria being met, I am interested in what a philosophical judgment “looks” like and how much the ideal story of data collection is corroborated by subjects’ verbal protocols. Such were open questions going into my experimentation: Will subjects think through their decisions at all? Do they tend to second-guess? Will they be able to verbalize precisely what they think as they arrive at a conclusion? In short, I entered into this research based on the view that it is plausible that the think aloud method would unearth useful information about how people form judgments and make decisions based on them. The method appeared uniquely poised to reveal what is motivating subjects to select the responses they select on x-phi research surveys since it provides a record of the immediate thought processes

that subjects have as they interact with the philosophical content of a thought experiment without interfering (to a significant degree) with the thought processes that interest x-phi researchers.

In addition, I paired the think aloud method with the qualitative interview. What makes the qualitative interview so useful is that it gives subjects an opportunity to reflect on their experience, explain any confusing verbalizations from the think aloud portion of the experiment, expand on their thought process in hindsight, and generally clarify the process of developing their judgments and selecting their answers. With the interview, they can do all of this in an open-ended format. In addition, I used the interview time to ask a series of questions that were designed to elicit more elaborate and nuanced answers than those available on the quantitative survey they first responded to.

Interestingly, my combination of the think aloud method and the qualitative interview is similar to what van Someren et al. describe as “thinking aloud – retrospection,” which combines the think aloud method with retrospection:

The think aloud protocols or behavioural observations during a session are used to obtain a retrospective protocol on pauses in the think aloud session or on fragments of the think aloud session that sounded incomprehensible, very incomplete or very odd. If possible this should be done directly after the think aloud session. If this is not feasible, one may show the written protocol to the subject and discuss the pauses, incomprehensible parts, etc. in a later stage.²⁴²

Here, the recommended supplement to the original think aloud session is a retrospective protocol that allows the subject to clarify specific moments within the original session. My research, in combining the think aloud session with the interview, follows this same structure, but treated the retrospective part as a full interview that included some questions that allow for clarification as

well as others aimed at revealing useful information about the think aloud session, the contents of the philosophical ideas being studied, and so forth.

Before conducting my study, it was my belief that the combination of the think aloud method and the qualitative interview would provide a body of information that is directly related to the ideal story of data collection. It would likely reveal just where this story is coherent and cohesive, and just where it is disjointed. It was my hope that the aggregate of the quantitative survey results, the qualitative protocols, and the content from the interview would uncover more precisely what goes on at the level of conscious awareness when one reacts to a thought experiment, forms a judgment, and responds to a question posed on the basis of the content within the thought experiment. It seemed that the combination of these methodologies is uniquely poised to reveal this information as each method directly and unobtrusively allows researchers to “hear” the thoughts and their impacts on a person’s cognitive engagement with a philosophical idea.

Though I wasn’t thinking about my project in terms of triangulation when I first conceived of it, it clearly speaks to that practice: the philosophical judgment, being the object of study, was to be observed from multiple methodological angles, including the survey, the think aloud method, and the follow-up interview. By using these methods together, I aimed to see a more total picture of laypeople’s philosophical judgments. In my mind when designing my study, this approach to research aligned very much with the x-phi research program even if it is a divergence in terms of methodological practice. That said, I still wanted to find a genuinely open place for my study to fit in with the movement that I took myself to be joining. For that reason, I explored the question, presented in the following section, as to whether x-phi has a place for qualitative methodologies and other new methods.

IV. Is Experimental Philosophy Open to New Methodologies?

Philosophical movements, rightly labeled, carry momentum. Wanting to join the momentum of the x-phi movement, and present my work within the most meaningful context available, I now make the case that x-phi is open to qualitative methodologies and that my study can be rightly classified as “experimental philosophy.” In short, I want to couch my research in terms of being an experiment *in* experimental philosophy because I think this labeling can help break open new possibilities for future research in x-phi.

My own research has the potential to serve as a strong criticism of x-phi research, and thus might be viewed as contrary to the movement’s vision. However, my study is grounded on a principle that serves as the foundation for all x-phi research: empirical observation is relevant to understanding philosophical ideas and the ways people think about them. In this sense, it might appear as if I am ironically employing empirical observation—via the think aloud method and qualitative interviews—to argue for the problematic use of empirical research methods in conjunction with philosophical thinking. However, this is not the case. I believe that empirical observation is relevant to understanding philosophical ideas and philosophical thinking, and I am motivated to scrutinize certain x-phi methodologies by making use of others. I am not focused on discarding x-phi research methods, but refining and rethinking them so as to better understand how people think about ethics and more. It is in this sense that my research is an experiment *in* experimental philosophy; I am using empirical research methods to investigate other empirical research methods to better understand the efficacy of the latter through the former. I am also interested in learning about people’s philosophical judgments/intuitions through new methods and through triangulation of methods; so in this sense my goals are very aligned with the x-phi movement.

Ultimately, I view my present research as a contribution to ethics as much as it is one to x-phi proper. However, experimental philosophers and others might not be willing to categorize the qualitative methodologies I employed as “experimental philosophy.” While finding the right category for my research is not the main focus of this project, it is important that I argue for bearing the label “experimental philosophy” for a few reasons. First, by being viewed as working in accord with the stated goals of experimental philosophers, my research might be taken more seriously by the people I want to take note of it, experimental philosophers in addition to ethicists. Second, bearing the label ideally allows experimental philosophers to trust that I am familiar with x-phi research and that my work is motivated by a desire to progress the movement, not impede it. Lastly, if qualitative methods such as the think aloud method and the qualitative interview prove useful and are categorized as within the bounds of x-phi, then it is plausible that more experimental philosophers will consider using them. In some abstract sense, labels *do not* matter: good and relevant research is good and relevant research regardless of disciplinary categorization. Practically speaking, they matter quite a bit. Whatever is considered the “standard” set of tools used by experimental philosophers will likely influence the design and execution of future x-phi research.

Experimental philosophy is sometimes defined or summarized both as a general, abstract notion of what experimental philosophy is *and* as a more concrete description of what experimental philosophers do. For example, the experimental philosophy topic page on PhilPapers opens with the definition that experimental philosophy “is a new movement that uses systematic experimental studies to shed light on philosophical issues.”²⁴³ This, in my view, paints experimental philosophy abstractly as a movement that encompasses a broad range of empirically informed philosophical inquiries—i.e., it articulates what experimental philosophy is, or could be, in theory. The next

sentence, then, brings experimental philosophy to something more concrete by claiming that “experimental philosophers apply the methods commonly associated with psychology...to address the kinds of questions that have been traditionally associated with philosophy.”²⁴⁴ Here, experimental philosophy is associated with the common practices that experimental philosophers use. This sentiment is further voiced in the statement that follows from the previous quote, which articulates that “a shared methodology” more so holds x-phi together than some “shared research agenda or metaphilosophical viewpoint.”²⁴⁵ On this definition, it is specific methodological practices common to experimental philosophy research that seem to define the movement more so than an abstract notion of using systematic experimentation to inform philosophical investigation.

However, not all conceptions of experimental philosophy are as restrictive. Experimental philosophers Joshua Knobe and Shaun Nichols, for example, open their piece “An Experimental Philosophy Manifesto” by aligning experimental philosophy with a traditional picture of philosophy in which “philosophers thought about human beings and how their minds worked. They took an interest in reason and passion, culture and innate ideas, the origins of people’s moral and religious beliefs. On this traditional conception, it wasn’t particularly important to keep philosophy clearly distinct from psychology, history, or political science.”²⁴⁶ Nothing in this depiction points to strict adherence to specific experimental methodologies, and nothing seems to define the experimental philosopher as a psychologist interested in philosophical ideas. Much like the PhilPapers description of experimental philosophy, Knobe and Nichols *do* move from this abstract notion to something more concrete—e.g., they write that experimental philosophers “proceed by conducting experimental investigations of the psychological processes underlying people’s intuitions about central philosophical issues”²⁴⁷—but their vision is still much less restrictive. In fact, they conclude their manifesto by claiming “the thing to do now is just to cast

off our methodological chains and go after the important questions with everything we've got.”²⁴⁸ Here, Knobe and Nichols are advocating for abandoning methodological constraints rather than being defined by them.

It was in this vein of discarding methodological restrictions, of viewing experimental philosophy as an investigation of humanity in all its complexity, that I designed my study. In particular, I want to learn about how people reason on issues in ethics with whatever tools will help. Specifically, I wanted to try a new approach to experimental philosophy research by making use of the think aloud method and the qualitative interview. From my research of the experimental philosophy literature, experimental philosophers had not employed either method of collecting qualitative data. And, more specifically, experimental philosophers had not explored the potential benefits of combining these two methods with the quantitative methods of data collection common to x-phi research. Without trying these two qualitative methods in conjunction with quantitative research methods, we would simply not know if the combination could help experimental philosophers with their varied goals. Therefore, to see what good the think aloud method and qualitative interviews might bring to understanding ethics, I conducted my study and analyzed the results as they pertain to a specific topic of interest, namely the principle explored by Chituc et al. that “ought implies can.” Ultimately, my research is a kind of experiment *in* experimental philosophy. Prior to having my data in hand, it was the case that my work may have revealed unique benefits to experimental philosophy from collecting qualitative data via these qualitative methods, or it might not have done so. Either way, this study was designed such that it would shed some light on the prospect of placing the think aloud method and qualitative interviews into the experimental philosopher's methodological toolbox.

In my view, the best way to put these methods to the test was to use them in conjunction with previously established quantitative methods of x-phi research. I am in agreement, then, with Andow who thinks “the most important contribution [qualitative methods] have to make is in supplementing the methods already used by experimental philosophers.”²⁴⁹ Therefore, rather than rely solely on quantitative or qualitative data—each which gives rise to unique benefits and drawbacks—I think it best to employ both kinds of methodology. This will potentially allow for each methodology to provide unique data that can shed light on data collected via the *other* methodology. This is the crux of triangulation of methods. This combination of qualitative and quantitative methodology might help redefine the tools that philosophers think of as useful for examining the world of moral philosophy and beyond.

CHAPTER FIVE

Design and Quantitative Results of My Study:

Adding the Think Aloud Method and Interview to

“Ought Implies Can”

I. OIC Experiment 3, My Study, and their Connections to the MPEP

If the reader is showing up in this chapter to get right to the heart of this dissertation, let me get her up to speed. In the interest of learning more about people’s philosophical judgments/intuitions on issues in ethics and beyond, I completed this dissertation as an experimental examination of the very process of capturing judgments/intuitions within x-phi studies, a process that often relies on quantitative methods of data collection.²⁵⁰ I used two related *qualitative* methodologies, the “think aloud method”—marked by having subjects talk through their thinking while engaging with a specific task²⁵¹—and the qualitative interview—marked by open-ended yet targeted questions presented to the subject regarding a specific experience.²⁵² I was motivated by the belief that each might have a unique capacity for shedding light on exactly what occurs during the experimental process of prompting and collecting a person’s philosophical judgments/intuitions. Specifically, I thought that applying these two qualitative methods to the experimental philosophy data collection process in the form of a unique experimental investigation would reveal that the quantitative data that experimental philosophers often collect and lean on is not reliably reflective of subjects’ judgments/intuitions regarding a given philosophical issue. I thought that these two methods, when applied during an empirical examination of x-phi methods, would show that the data being collected in x-phi research via quantitative surveys cannot be

appropriately and uniformly labeled as “judgments” or “intuitions” with reference to the specific philosophical issue that is presumably being examined through experimentation. This project is especially relevant to moral philosophy, because in my experience people have rich and nuanced judgments on issues in ethics that are not easily captured with quantitative questions alone.

Let me translate this into terminology I’ve crafted and used thus far in the dissertation: I was motivated by the *methodological problem of experimental philosophy (MPEP)*, a problem borne out of the lack of epistemic justification experimental philosophers have in thinking that the quantitative data they’ve collected based on surveys are reflective of participants’ genuine, relevant philosophical judgments toward a philosophical question (or questions), such as those in ethics. To evidence my worry about the MPEP, I discussed it earlier in various ways in relation to multiple x-phi experiments. My most fine-grained version of this discussion highlighted the breaking down of the *ideal story of data collection*, the assumed way that philosophical judgments are elicited and captured by quantitative surveys, in relation to a study on moral obligation by Chituc et al. that focused on the specific Kantian principle that “ought implies can” (OIC).

With my motivating concerns and research ideas in mind, I went forward with my study, which focused on a specific experiment conducted within a three-experiment study by Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard, described in their paper entitled “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can.’”²⁵³ I chose this study (and the specific experiment within) for various reasons, which I highlighted in chapter three, but in large part I did so because it is representative of the kind of x-phi research I am interested in analyzing. My experimental research contained two main parts. First, I conducted an in-person replication of part of the third experiment (OIC experiment 3) within this three-experiment study by Chituc et al. (the

OIC study), which was originally conducted using Amazon Mechanical Turk, though my version was modified in that I applied the think aloud method during subjects' survey completion (along some other small changes I'll discuss below). Then, in the second part, I conducted qualitative interviews with subjects which focused on their experiences completing their surveys and thinking aloud. Together, these two main parts allowed me to collect qualitative data *on the data collection process* used by Chituc et al. in OIC experiment 3, an experiment that is representative of many x-phi studies.

While I didn't have all the details of the term "MPEP" hammered out when I designed my study as a modified version of the OIC study, the MPEP was in essence what I sought to explore. I thought there really was a methodological problem such that experimental philosophers couldn't justify their findings in a compelling enough sense. And I really thought I would reveal this most compellingly by showing that they weren't actually justified since their data wouldn't turn out to be reflective of people's relevant philosophical judgments/intuitions. Beyond that general prediction, I wasn't exactly sure what I would find in the qualitative data until I recorded and analyzed it. So let's move closer to that data analysis by getting to the specific details of my research design.

II. Experimental Design

So, what is the experiment (OIC experiment 3) that I ended up modifying? Taking their first two experiments to support H2, which is roughly a negation of OIC, rather than H1, which is roughly an affirmation of OIC, Chituc et al. further explore the relationships between "ought" and "blame" and "can" in experiment 3.²⁵⁴ Specifically, they design a study that addresses questions both that I was and wasn't interested in, which explains why I only partially recreated their work.

Of interest to me were two of eight related thought experiments that pit “ought” against “can” in an interesting way: participants are asked about an agent who can no longer make a promised movie date with a friend because he either suffered car trouble or had a change of heart.²⁵⁵ Not of interest to me were the other six related thought experiments that didn’t pit “ought” against “can” but nonetheless allowed Chituc et al. to explore participants’ understanding of “non-moral” conceptions of obligation, depicted by cases wherein an agent intentionally or unintentionally fails to go to a movie by himself, as well as participants’ understanding of obligation and blame in connection to instances *before* an agent can no longer carry out a plan.²⁵⁶

For the two scenarios I used for my study, which I will detail below, Chituc et al. asked participants if they agreed or disagreed with three statements based on a 100-point scale described in their first experiment (with -50 at one extreme, 0 as neutral, and 50 at the other extreme).²⁵⁷ The three statements asked if Brown, the promisor, 1) was able to make his meeting with Adams, 2) was blameworthy for failing to make the meeting, 3) ought to make his meeting.²⁵⁸ In my study, these same three questions, though with different but mathematically proportional scales, were always presented in this same order (as presented in the example in their publication²⁵⁹): can, blame, and ought. However, there is no indication if Chituc et al. administered these three questions in random order or fixed order. If it turns out that Chituc et al. did present their subjects these questions in random order in OIC experiment 3, then that is going to be another potentially relevant difference between my study and theirs.

I was ultimately less interested in examining the specific claims made by Chituc et al. in the framing and analysis of their results from experiment 3, and more interested in seeing if the data they produced seem to reflect genuine, relevant philosophical judgments. While I will, here and there, engage with the general conclusions Chituc et al. make—that their work bears on the

relationship between “ought” and “can”—I am less concerned in grappling with the specific ways they interpret *this* or *that* data analysis to support or negate H1 or H2. I don’t want too much of my work to be targeted at these hypotheses when it isn’t clear to me that 1) H1 and H2 are adequate for framing the experiments from Chituc et al. insofar as we are talking about the Kantian relationship between “ought” and “can,” and 2) the data collected are reflective enough of genuine, relevant judgments so as to even speak in support of or in opposition to H1 and H2. Nonetheless, it is helpful to see what Chituc et al. make of their own data from OIC experiment 3 in conjunction with the other experiments. In large part their conclusion, as it deals with moral obligation, adds up to support for H2 and a negation of OIC for the reason that people’s judgments of “ought” were influenced by “blame” rather than tracking with “ability,” which undermines the entailment between “ought” and “can.”²⁶⁰ And here’s where I want to jump in with my study, because these claims can only be meaningful as such if the ideal story of data collection holds up; they can only be taken as established if the researchers don’t run into the MPEP. We’ve already seen there is good reason to think the ideal story doesn’t hold up in experiments 1 and 2 based on the limited qualitative data we had access to, and it is likely the same is true of experiment 3. It is likely that Chituc et al. are facing the MPEP squarely with OIC experiment 3. While there was no qualitative feature of OIC experiment 3 in the OIC study, my own study presents a great deal of such data in the form of think aloud protocols and qualitative interview transcripts which will reveal just how deep the MPEP runs with relation to OIC experiment 3.

What follows in this section, now, is a breakdown of my experimental design in terms of relevant sub-sections. From there I present my quantitative results, and the next chapter begins outlining my qualitative data.

Recruiting Subjects

Chituc et al. recruited 321 subjects via Amazon Mechanical Turk and paid each subject 30 cents for his or her completion of the experiment²⁶¹—in total, though, there were just 84 participants (split 48 and 36) in their study who completed the two conditions I replicated.²⁶² Because I was interested in collecting extensive qualitative data rather than establishing statistically significant effects based on quantitative questions, I was not in need of recruiting as many subjects. I worked with a total of 40 subjects: 20 of whom read about Brown’s *high fault* (conscious decision) inability to carry out his moral obligation to Adams, 20 of whom read about Brown’s *low fault* (car troubles) inability to carry out his moral obligation to Adams. And because of my strong preference for conducting this study in person, and the fact that think aloud and interview methodologies translate best to in-person interactions, I worked with all participants in person at Claremont Graduate University’s Philosophy Faculty Offices building. Also, I wanted to pay subjects significantly more, \$25, for making their way to the campus building, completing one or two warm-up surveys while thinking aloud, completing the main survey while talking aloud and the interview while I recorded the audio throughout, and completing a brief survey following their completion of the experimental portion which included demographic information and questions about their level of familiarity with Western analytic philosophy as well as the specific debate around “ought implies can.” The heart of my modified version—the two quantitative surveys about high and low fault moral failures—is the same as Chituc et al. used, so I expected my quantitative results (explored in the next section of this chapter) from my 40 participants to be similar to the same conditions in OIC experiment 3. Overall, they were similar.

I recruited subjects using a variety of methods. Once a subject was recruited, he or she met me at the Philosophy Faculty Offices house at Claremont Graduate University to complete the study. All subjects were recruited in one of four following ways:

1. Through recruitment flyers I designed (flyer is available as Appendix F) which contained my email for setting up an appointment;
2. Through a standardized email that I designed (email body is available as Appendix G) which contained my email for setting up an appointment;
3. Through word of mouth, by asking participants who completed the study if they knew anyone who might be interested in emailing me to set up an appointment;
4. Through word of mouth by asking people I know to participate.

Though these the four methods account for the recruitment of all 40 subjects, I don't specify how many subjects were recruited by any given means. For any subject that I recruited from people that I knew prior to interacting with them for the study, I was confident that he or she had no prior knowledge of my research on "ought implies can" or what the content in the vignettes would be about. Nonetheless, there is surely something comforting in using a tool like Amazon Mechanical Turk in recruitment because one can be nearly certain that some online subject would have never encountered your project in virtue of his or her knowing you in some capacity. And perhaps the worry about potential bias is even more pronounced for those subjects recruited by previous subjects: might some subjects have recommended a friend and given them some quick details about my study? It's certainly possible. But similar worries crop up for studies that only send out flyers or emails for recruitment: whether researchers *ask* for subjects to recruit other subjects or not, subjects can point their friends to a flyer, or an email, etc. and suggest they sign up and, just the same, provide some information about the study itself. And, it isn't as if knowing that a survey

was sent to a large pool of people on the internet secures us from all potential biases or interferences: while I can't be certain that participants who passed on my email to someone didn't tell them a bit about the content of the study, I can be sure that subjects didn't use the internet to look up the definition of "ought" (or more problematically "ought" and "can") while completing my study, something users of Amazon Mechanical Turk might very well have done.

While I think these possible flaws and biases should be addressed, they really aren't worrisome in my view. Even if someone somehow knew my views on OIC prior to completing the study or if they looked up "ought implies can" online to learn what the accepted philosophical view is, I just don't know that it would have influenced their long-form interview/discussion with me about the topic. People just don't seem to have so much at stake in whether "ought" implies "can." If there are biases worth worrying about, I think they are far more common and difficult to eradicate. For example, I suspect many people participate in experimental studies with a bit of misconceived suspicion that there is some trick or manipulation that is embedded in the experiment *even if* they are told straightforwardly that this isn't the case. I'm much more worried about subjects avoiding straightforward responses because they think there must be some *hidden* meaning that the researcher is really after, even if there isn't. And I'm definitely much more worried that the very questions I constructed for my interviews prime subjects for *this* kind of response over *that* one—i.e., I'm more worried about biases seeping into the questions I wrote, the order I presented them, and the way I spoke them than I am any inter-subject collusion. I think other x-phi researchers should be, too. The best way to deal with such biases is for others to read my interview excerpts and criticize instances in which I may have acted on bias *and* for future researchers to conduct similar interviews to see if similar responses result.

The Scripts

I greeted subjects when they arrived at the front door of Philosophy Faculty Offices house, and then I brought them around to the conference room. When they were ready to begin, I started reading my script for the study at the top, telling them I would give them an informed consent form to read and sign. Throughout the study, I used two versions of the script (available as Appendices I and J) that were subtly different. The reason I updated the first one (available as Appendix I) to make the second one (available as Appendix J) was to clarify some information and emphasize *what* I was interested in studying so that subjects weren't worried about some hidden purpose. For example, amongst other small changes, I added the following line to the second version of the script: "As mentioned on the informed consent form, I am interested to learn how people think about common philosophical ideas." It's also worth noting that while I always conveyed the information from my script, I sometimes spoke the information in a somewhat conversational manner such that I would look up at the subject and complete a sentence in its essential content from memory rather than glue my eyes to the page and read the script verbatim in a monotone, overly-impersonal way. I might have elaborated on certain sentences, paused after others, etc. After each part of the study, I would again return to the script to get things moving to the next part. So, after administering the informed consent form, it was time to administer the warm-up surveys. I used the first script for the first four participants, and the second one for the remaining 36.

Warm-Up Surveys

Based on the script, I next moved to have subjects warm-up in thinking aloud. The main goal with incorporating this method in the study was to capture the qualitative data *over and on*

top of the quantitative results so that I could collect any relevant information that enables me to better understand the ideal story of data collection within x-phi studies of this kind. Specifically, from my script, I read the following: “When I say ‘begin,’ please start reading aloud from the top of this page and please make sure to say out loud any thoughts you have while reading and completing the survey. In short, you are being asked to ‘think aloud’ and talk your way through this survey. I will not be recording the audio from this Warm-Up exercise unless you request that I do so.” This meant that subjects were to read the vignette aloud, the questions aloud, and their responses aloud, and they were also to verbalize any thoughts they had before, during, and after this process.

I got the idea to incorporate warm-up surveys from van Someren et al.’s book *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes*.²⁶³ I gave one subject (participant 1) just one warm-up survey (the one entitled “Lying” which is detailed below), and I gave the rest of the subjects two warm-up surveys (including “Lying” and the other one entitled “Giving,” also detailed below). In discussing warm-up sessions, van Someren et al. suggest using a practice task similar to the main one.²⁶⁴ Heeding this advice, I designed and administered two surveys that were somewhat similar to the ones from the OIC study, but not too close to their philosophical content. My warm-up surveys each involved agents dealing with mundane moral situations, but neither warm-up dealt with the relationship between “ought” and “can.” Though I had labeled the warm-up surveys 1 (“Lying”) and 2 (“Giving”), I didn’t record the order that they were administered to subjects so it is possible that some subjects read 2 before 1.²⁶⁵

The first warm-up survey I created was entitled “Lying” and it included the following vignette (the full survey is available as Appendix A²⁶⁶):

Chris has told Amy that it is always wrong to lie, even if the lie will prevent something bad from happening. One day, Chris lies to his employer and says that he is sick and cannot attend work. Chris lies so that he can spend the day donating his time at the local retirement home talking with and caring for seniors who usually get few visitors.

Chris tells Amy about the situation and admits that he did in fact lie and that he did in fact use the day to help out at the local retirement home.

In response, subjects were asked the same kind of question that subjects in OIC experiment 3 were asked, whether or not they agree with a particular statement. However, unlike the 100-point scale used by Chituc et al., I used a 10-point scale that was proportional to their 100-point version. Since their surveys were completed on a computer (on Amazon Mechanical Turk), I imagine it was easier to create a 100-point scale that subjects could navigate. On paper, the 10-point version maps on better, which is why I opted for -5 and 5 to represent the extremes rather than -50 and 50 (see appendix A for full survey including my scale). Specifically, my subjects were asked if they agreed or disagreed with these two statements on this first warm-up survey (“Lying”): 1) “Amy has good reason to think Chris is a frequent liar.”; and 2) “Amy has good reason to distrust Chris in the future.” I did not record audio while subjects completed this warm-up survey or the next one unless they asked me to do so (per the researcher script) but I did ask them to think aloud as they did so.

The second warm-up survey I created was entitled “Giving” and it first included the following vignette (the full survey is available as Appendix B):

For the last ten years, Josephine has donated 5% of every monthly paycheck she earned to a local charitable organization. This month, however, Josephine decides to spend the 5% percent of her paycheck on a new flat-screen TV.

Brad, who manages the local charitable organization, noticed that Josephine’s check did not arrive—very few people donate money so it was obvious when Josephine did not do so.

In response, subjects were asked if they agreed or disagreed with the following three questions (based on my 10-point version of the Chituc et al. scale): 1) “Josephine should have called the local charitable organization to let them know that she would not be donating this month.”; 2) “Brad has every right to be frustrated that Josephine did not donate this month.”; and 3) “Brad has every right to be frustrated that Josephine didn’t call the local charitable organization to say that she wouldn’t be donating this month.” Once the warm-up session was completed, I moved on in my script to the main surveys during which I would record subjects as they thought aloud.

The Main Surveys

After the warm-up surveys, I moved forward on my script to introduce the main survey that a given participant would take. 20 of the 40 subjects completed my version of the Chituc et al. survey in which an agent failed to complete a moral obligation because he decided to wait too long so he couldn’t make it on time. I refer to this survey as Main survey 1, “Moral Obligation High Fault Unable” and it is available as Appendix C. Here is the vignette, which is identical to the one from Chituc et al.:

Brown is excited about a new movie that is playing at the cinema across town. He hasn’t had a chance to see it, but the latest showing is at 6 o’clock that evening. Brown's friend, Adams, asks Brown to see the movie with him, and Brown promises to meet Adams there.

It takes Brown fifteen minutes to drive to the cinema, park, purchase a ticket, and enter the movie. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time.

As Brown gets ready to leave at 5:45, he decides he really doesn't want to see the movie after all. He passes the time for five minutes, so that he will be unable to make it to the cinema on time. Because Brown decided to wait, Brown can't meet his friend Adams at the movie by 6.²⁶⁷

And for the follow-up questions from my version of the survey, I always asked the same three questions in the same order. Participants were asked if they agree or disagree, based on a scale from -5 to 5 rather than Chituc et al.'s scale from -50 to 50, with these statements: 1) "At 5:50, Brown can't make it to the theater by 6."; 2) "Brown is to blame for not making it to the theater by 6."; 3) "Brown ought to make it to the theater by 6."²⁶⁸

The other 20 completed my version of the Chituc et al. survey in which the agent failed to meet a moral obligation because of car troubles. I refer to this one as Main survey 2, "Moral Obligation Low Fault Unable" and it is available as Appendix D. Here is the vignette (which is originally from Chituc et al.'s study):

Brown is excited about a new movie that is playing at the cinema across town. He hasn't had a chance to see it, but the last showing is at 6 o'clock that evening. Brown's friend, Adams, asks Brown to see the movie with him, and Brown promises to meet Adams there. It takes Brown fifteen minutes to drive to the cinema, park, purchase a ticket, and enter the movie. It would take 30 minutes if Brown decided to ride his bike. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time.

At 5:30, Brown thinks about riding his bike, but decides it is too cold. Instead, he leaves at 5:45, but his car breaks down five minutes later. He can't fix it himself in time to make it to the cinema, and it is too late to make it by bike. Because his car is not working at the time, Brown can't meet his friend Adams at the movie by 6.²⁶⁹

And the same questions listed above were used in conjunction with this vignette.

At this point I want to note an important difference between my version of the surveys and those from Chituc et al. In both of my surveys (Appendices C and D), I used reverse language on my question about whether or not Brown *can* make the appointment to the same effect as Chituc et al.'s question—i.e., subjects in my study and theirs both reported as to whether they think Brown can make it—but with inverse responses on the scale since subjects would *disagree* (a negative numeral) with Brown's *ability* in their study and *agree* (a positive numeral) with Brown's *inability* in my study. I misunderstood how Chituc et al. phrased the ability question in their conditions because I misinterpreted their writing. For my study, then, subjects *agreed* that Brown can't make his meeting, which isn't so unnatural, but then if they thought Brown could make it, they'd have to disagree that he can't—thus creating an awkward double negative. This, though, explains my positive score to their negative on just this question (see Figures 2 and 3 in the next section): with can't, my subjects *agreed* (a positive number on the scale) to affirm Brown's inability; with can, their subjects *disagreed* (a negative number on the scale) to affirm Brown's same inability. In the end, we have the same general result in that Brown *can't* make it. This small error turns out to be less than troubling in that my study involved recording qualitative data in the form of a think aloud transcript and a follow-up interview, so I can see if subjects were getting things backwards or getting confused with the scale. For example, the interview transcript made it clear that participant 12 ended up switching her original response to this question once she realized that she meant the

opposite of what she initially selected.²⁷⁰ While this isn't ideal, I have in some sense demonstrated a clear value of my qualitative methods just in that my qualitative data helped me see if someone had incorrectly understood the question she was being asked. Without the transcript data on hand, I would be forced to assume that *each* participant selected what he or she selected based on his or her understanding of what was being asked. And there are likely many questions in quantitative studies that don't read as clearly as they are intended to read, and in those cases having access to interview data about the surveys can bring to light any confusions or misunderstandings.

To summarize all of this, the two surveys I used represent two of the eight surveys that subjects in OIC experiment 3 received but with some important changes: my subjects always got the questions in the same order (which I'm unclear as to whether that is true of OIC experiment 3 subjects), my subjects completed questions on scales from -5 to 5 rather than -50 to 50, my subjects were asked to agree that Brown can't make it rather than disagree that he can, and my subjects completed the surveys in person *while* verbalizing their thoughts into a recording device rather than on a computer without thinking aloud. The two conditions were chosen because they most obviously pit "ought" against "can" in that the agent *can't* carry out his obligation in both surveys, he has or had an obligation to do so, and the key variable being altered is whether that inability comes from something external to his decision-making (car troubles) or internal to it (lack of want).

The Interview

After subjects completed one of the two main surveys, I continued in my script by saying that the interview would follow. I want to emphasize that subjects were encouraged to look at the survey they had just completed in response to the interview questions (see Appendices I and J). There is an obvious worry here that subjects might just look at their responses during the interview

and simply rationalize the choices they made rather than respond in earnest to my then present interview questioning. Had I prohibited them to look at their responses, this worry says, I might have ensured that subjects respond honestly to the interview, perhaps exposing inconsistencies between their initial survey responses and their follow-up responses during the interviews. This is a legitimate worry, and future studies might try removing the completed survey from the scene during the interview. However, I think this runs a different risk, one that I thought greater and more likely which is why I opted to keep the survey present: without their responses in front of them, subjects would avoid remarking on their survey completion in a direct way so as to avoid being inconsistent. For example, a subject asked about their response to question 3 might dance around saying what they think about the word “ought” in the question in case they forgot what they put and didn’t want to say something in contradiction. In addition, by having their surveys in front of them, subjects can react in real-time and express if they feel their responses might change upon further reflection, which wouldn’t be possible if they didn’t have their responses in front of them. That is, subjects might re-read and re-consider question 3 during the interview and remark that they chose x as their response because y but now they think z might more accurately represent their thinking. By having control of their initial responses, *they* get to decide if they want to revise them during an interview, but without access to their initial responses, they fail to have control in some sense and might be reluctant to offer revisions. This speculation might be very misguided, though, and it might turn out that having one’s answers in front of oneself is an invitation to provide inaccurate rationalizations rather than genuine attempts at explanation. That said, I tend to worry less about getting these rationalizations or justifications regarding such low-stakes, mundane thought experiments. Still, my decision here can certainly be questioned or reconsidered.

So, with surveys in hand, participants were interviewed in a qualitative interview process that included questions about their experiences completing the survey as well as other questions related to the content of the thought experiment and OIC more generally. During the interview, I recorded audio, which I later transcribed. My detailed transcription methods and their justification is available as Appendix M, and I analyze the participant transcripts after this chapter. During each interview, I had before me a set of questions that I largely stuck to, though I often jumped around in order, added follow-up questions, skipped questions, etc. as would be within the scope of the original interview questions and relevant to OIC. For the first four subjects, I used the first set of representative qualitative interview questions (available as Appendix K) and for the rest I used the second set of representative questions (available as Appendix L). I changed my questions for a number of reasons that can be best summed up as I revised the script based on my learning from the first four interviews. I felt that I had a better sense of what I wanted to ask, and how I wanted to ask it so as to get at the interesting philosophical questions surrounding the thought experiments and OIC.

After the interview, I asked subjects to complete a brief follow-up survey without any more audio recording or thinking aloud. The survey is available as Appendix E, and it includes a number of demographic questions and other questions about subjects' familiarity with Western analytic philosophy.²⁷¹ From there, I wrapped up the experiment: I thanked participants, answered any questions they had about anything, asked if they might know anyone who would be interested in emailing me for an appointment, etc. And with that, I had both the quantitative and qualitative data ready for analysis. In the next section, I'll analyze the quantitative results and their relation to those from OIC experiment 3; in the next chapters the focus will be on the qualitative data.

III. Quantitative Results and Analysis

In the next chapter, I get to what I most sought after, the qualitative results of my study. For now, I want to highlight my quantitative findings. The quantitative results from my study are reported in Appendix H which is a spreadsheet with two tabs (I'll primarily be referring to the first tab). Before getting to the core findings and my statistical analysis, I want to briefly address what I *won't* be making much of here in this dissertation.

Columns F through L are identical on both tabs of the data spreadsheet (Appendix H), and both contain data from the follow-up survey (Appendix E). In my qualitative analysis of the data in the next chapters, I will make some reference to information gathered from this survey as it is relevant. However, I won't make anything of the follow-up survey data in terms of statistical analysis. I have two main reasons for this: 1) my sample size of 40 subjects is quite low for the kinds of statistical analysis needed to make something of any trends in connection with subjects' gender differences, age differences, etc.; 2) I'm not interested in making claims about how certain of the follow-up survey data connect with subjects' responses in some generalizable fashion because my research is not geared toward generalization as much as it is concerned with exposing instances and patterns in subjects' qualitative responses that are worth paying attention to if we want to understand subjects' judgments toward OIC.

In fact, I really don't think I can or should make much of my quantitative data on the main survey questions either—that is, in addition to the follow-up survey, I don't think I have much to proclaim even on my collected quantitative data from the core questions derived from Chituc et al.'s surveys. This is again because of my low sample size (relative to quantitative survey research sample sizes). However, I *will* end up proclaiming a few things through my statistical analysis of the main survey question data because I wanted to see and did in fact find that I was getting results

similar to what Chituc et al. found. Did I perform something like a full-fledged replication? Not exactly. But I did find, as will be explored later in this section, that my results generally match those of Chituc et al.

Also, I want to qualify some of the data from this follow-up survey so that readers don't take more from it than seems appropriate. Simply put, I just don't think question 3 on this follow-up survey (available as Appendix E) is always a good measure of what I was hoping it would measure. Here is question 3: **"Prior to completing this study, were you familiar with the philosophical idea that *ought implies can*? Please circle your response."** In my mind, this question would capture the almost certainly true fact that few people, even those who have taken one or more philosophy classes, are familiar with the *philosophical* debate around OIC as codified by Kant. I designed this question to see if people had read or discussed *not merely* issues of obligation and ability, but instead the claim common in philosophy literature that ought entails can, or the pushback that perhaps it doesn't—to answer a 2 out of 5 would indicate something like you encountered the ideas of ought and can in some specific fashion in discussions of moral philosophy, and to put a 5 would be to indicate that you know OIC's connection to Kant and his justification for it, or perhaps you know the larger debate borne out of his proclamation or similar proclamations. And yet, I got the impression that some subjects answered in such a way that might have overrepresented their familiarity for one completely understandable reason or another. For those who did declare familiarity with OIC, perhaps they took their familiarity with the concepts of "ought" and "can" as indicative of their familiarity with the OIC debate in philosophy—something completely understandable given my unfortunately vague wording of the question. The question can be reasonably taken as a thermometer on whether someone understands how ability to act relates to obligation—nothing in the question gets narrow enough on the OIC *philosophical*

debate as I'd have liked. Or perhaps these respondents wanted to appear as if they were familiar, so as to be viewed as in the know. And, of course, some subjects might have been as familiar with the Kantian debate just as they reported. The upshot of all of this, then, isn't that none of the responses from question 3 should be taken seriously, but rather that there is a chance that some responses are inflated, so little substantial can be made from them. This is most obvious whenever a subject reported higher familiarity with OIC than concepts in philosophy (question 2) since such a knowledge-base is truly anomalous because OIC *is* a philosophical concept—participants 1 and 18 are such examples (see Appendix H).

One other thing to notice in the data points themselves before I jump to my statistical analysis is that there are two tabs on the spreadsheet (Appendix H) that contain a few different data. In terms of numerical data, the tabs are identical except for a few responses on the three main survey questions that a few subjects ended up changing during the interview *after* completing the survey. I recorded these changed responses, and I've marked them using the legends just below the data on each spreadsheet tab: the first tab, "Research Data, Orig. Data," has subjects' original survey responses; the second tab, "Research Data, w Switched Answ," has the data including the switched responses. In some cases, subjects changed their responses because they realized, only in the interview, that they misread or misunderstood a question. In other cases, they changed their responses based on some kind of reconsideration. I elaborate on these data further in Appendix N, and I also explain why I only ended up analyzing the data from the first tab.

For just one example of misreading or misunderstanding the question, participant 12 says the following in the interview when she is asked about her response to question 1 on the survey:

Umm..I said completely agree just because of the question it was asking...umm...it was saying he can't make it by oh 6 oh man I think I meant to say completely disagree sorry..[I

say “Oh OK”]. [unintelligible]. if I can change that. [I say “Yea that’s fine”]. . . OK [laughs while speaking] I just put it the wrong way [laughs]. . . umm. . . I thought it said c . . . I read can for some reason but anyway . . . umm. . . he can’t make it by theater by 6 and I said completely disagree because he can if he like got into another car.²⁷²

I think this example is especially important because it helps researchers think about how to design easy to follow survey questions—of course, this example is from a (previously mentioned) subject who misunderstood or misread the first question which was the result of my previously discussed error and so this says nothing regarding the question as written by Chituc et al. I think this participant can be forgiven for this error since she (along with all participants in my study based on OIC experiment 3) were asked to agree or disagree with a statement that agent “can’t” do something rather than a statement that agent “can.” It is just confusing to *disagree* that someone *can’t*, as this relies on a double negative to create an affirmative. Fortunately, though, the intended answer is that subjects agree that Brown can’t make it (since the vignette affirms this) so no double negative would arise for that “correct” response. And even if I hadn’t errored in flipping “can” and “can’t,” it is not necessarily intuitive to think as philosophers do and ask participants to evaluate a sentence as a whole as if it is a proposition that can be judged based on a scale of agreement and disagreement. Perhaps a better scale would place “Brown CAN make it” on one end and “Brown CAN’T make it” on the other end. Either way, it is helpful to have think aloud and interview transcripts to expose when subjects do read questions differently than we’d expect.

Lastly, before looking at statistical analysis, I want to reference Appendix N, which details some additional information, mostly trivial, about my coding choices and data analysis choices. In large part, it details how the data on Appendix H came to be recorded and coded as such in those moments when I had to make some kind of judgment as to how to code a particular value. For

example, I explain why I coded certain fill-in-the-blank responses on the follow-up survey one way or another. In an effort to be transparent, I have included this appendix that it might help someone understand how I came to record and report my data.

* * *

The first bit of data to explore are the means and standard deviations from OIC experiment 3 as well as my own data. To do so, I translated the data from OIC experiment 3 from scores in the -50 to 50 range to those in the -5 to 5 range. I simply divided each score from the two conditions that I partially replicated—the conditions for high fault/moral obligation/inability and low fault/moral obligation/inability. Then, excluding the same data points as did Chituc et al., from those two participants who failed the attention check in the low fault/moral obligation/inability condition, I calculated the mean and standard deviation for each response, which maps onto the data they produced in their publication:

**Figure 2: Mean & Standard Deviation from OIC Experiment 3:
Translated to My Scale of -5 to 5²⁷³**

		Can	Blame	Ought
High Fault	N	48	48	48
	Mean	-3.504	3.819	.617
	Std. Deviation	2.3142	2.6535	3.5358
Low Fault	N	34	34	34
	Mean	-3.474	-1.532	-3.006
	Std. Deviation	2.6938	3.0056	2.5605*

*This entry doesn't precisely correspond with that from OIC experiment 3; it appears they didn't round up on this value from the "5" in the 10,000th place for whatever reason.

Figure 3: Mean & Standard Deviation from My Study

		Can't	Blame	Ought
High Fault	N	20	20	20
	Mean	3.050*	4.650	2.350
	Std. Deviation	2.9643	.6708	3.0826
Low Fault	N	20	20	20
	Mean	3.700*	-.550	-1.300
	Std. Deviation	2.8116	3.6487	4.0275

*Both of these numbers, representative of the mean for the question of “can’t,” are positive values rather than negative values as would correspond with the data from OIC experiment 3. This is because I misunderstood the wording of Chituc et al.’s original question and asked subjects to agree that Brown can’t rather than disagree that he can. The same general result is exemplified in both their data and mine, but they are simply inverted in terms of positive and negative numerals as a result. Therefore, in my statistical analyses, I will invert my scores on this question so they are negative values and therefore correspond to the ones in the Chituc et al. data.

Next, I want to display my scores in conjunction with those from OIC experiment 3 in overlapping histograms so it is clear how closely my data did or didn’t sync up with theirs. Below, Figures 4 through 9 display those histograms.

Figure 4: High Fault Can Data Comparison

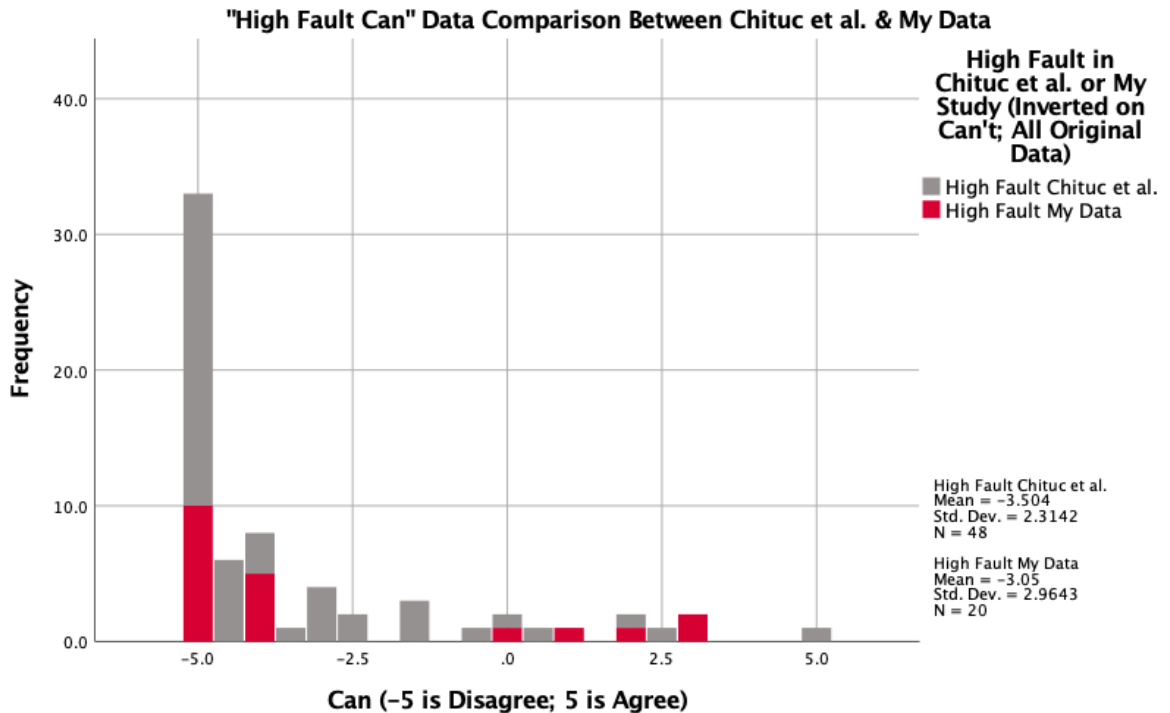


Figure 5: Low Fault Can Data Comparison

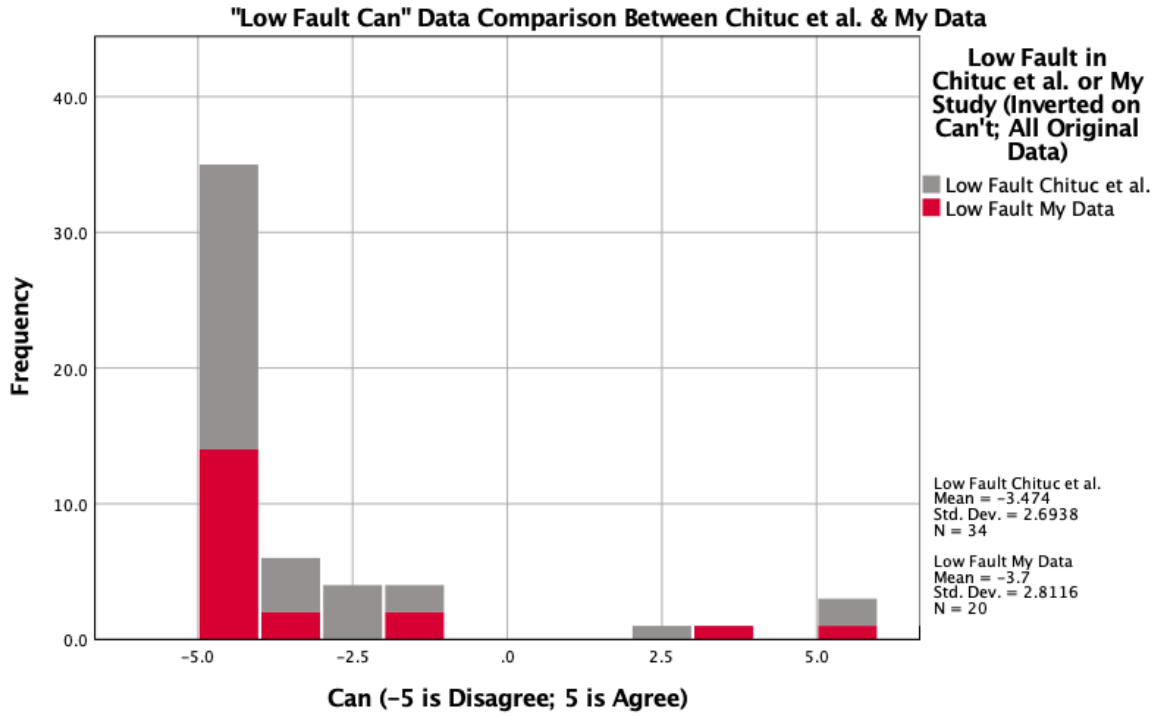


Figure 6: High Fault Blame Data Comparison

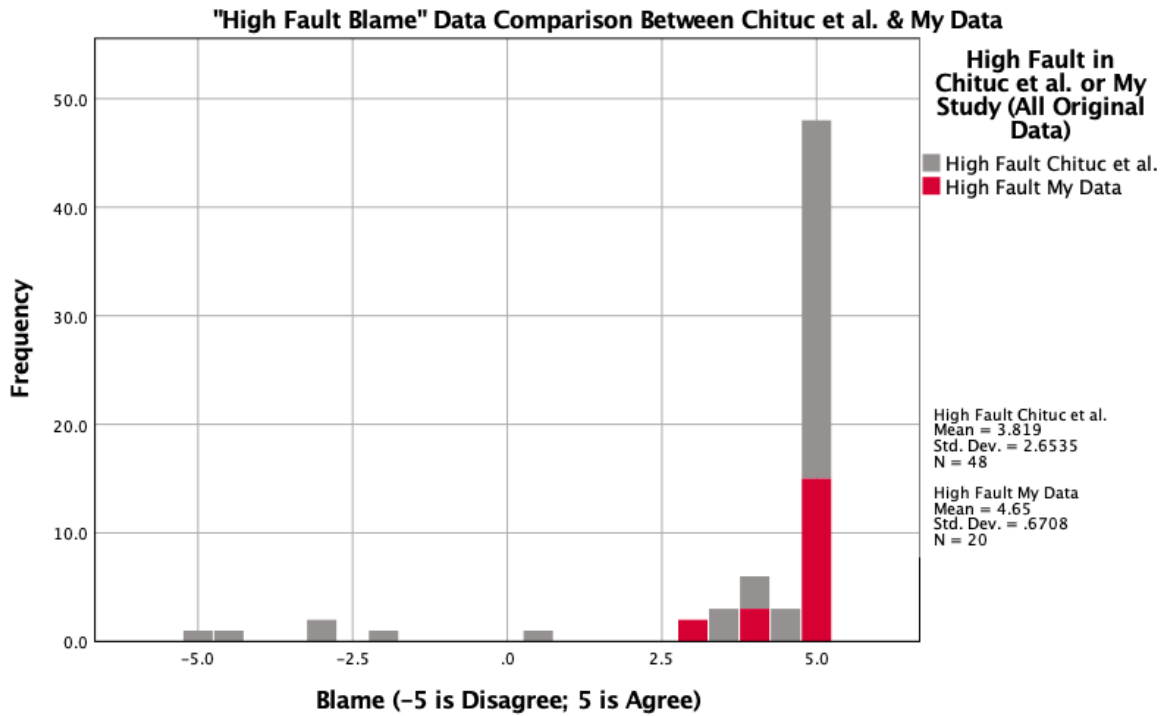


Figure 7: Low Fault Blame Data Comparison

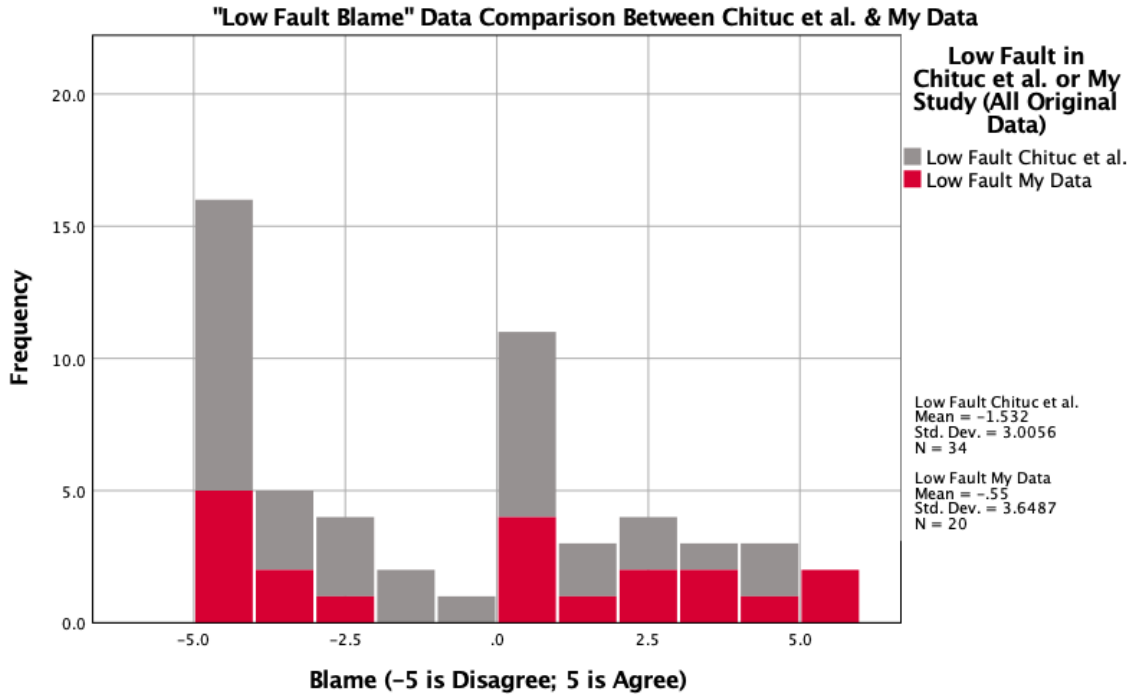


Figure 8: High Fault Ought Data Comparison

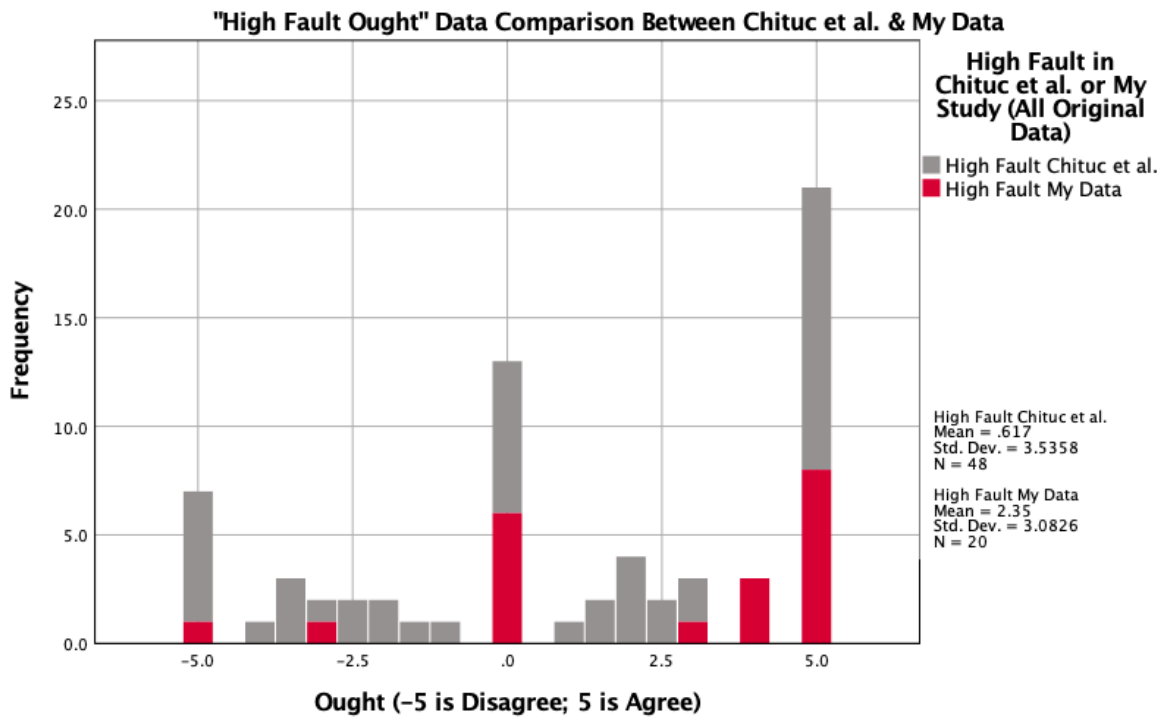
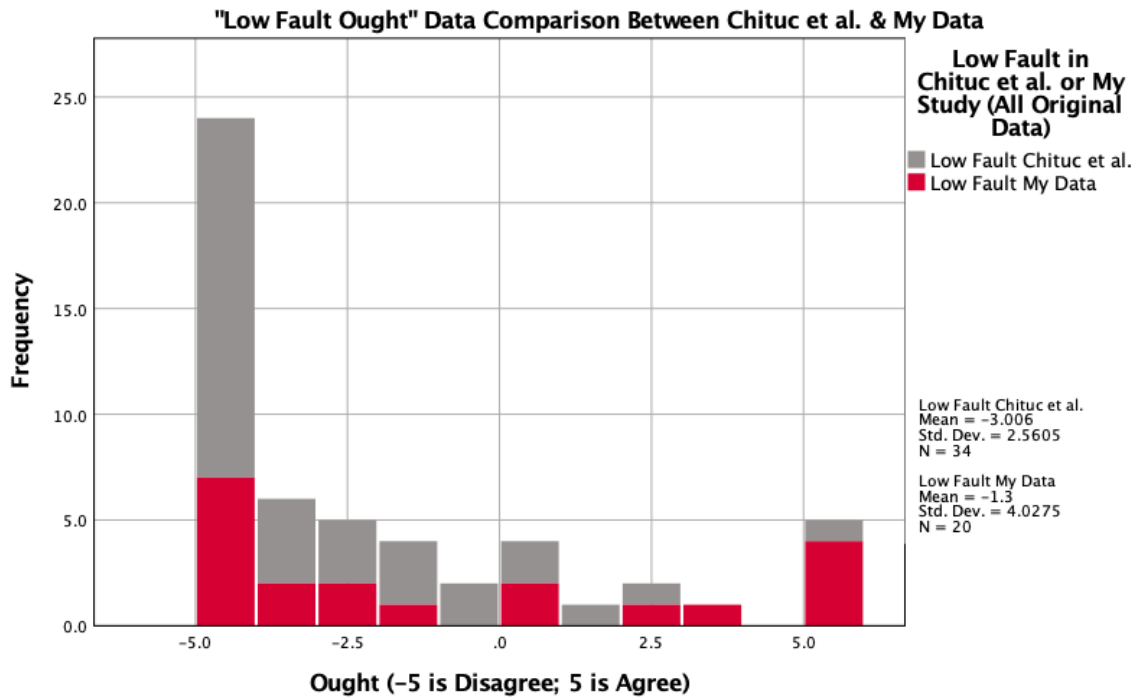


Figure 9: Low Fault Ought Data Comparison



As should be clear in these histograms, my data trends largely map onto those from Chituc et al. Of course, there are some differences and some nuances to explore (which I do below), but the general shape of my data syncs up with those from the study I partially re-created.

The next thing I want to look at is the correlations that Chituc et al. report just on those conditions in which Brown is unable to make the meeting and has a moral obligation to do so; these are precisely the conditions that I explored in my project. Although I punched in their data correctly, and produced identical means and standard deviations (as detailed above), I couldn't get the correlations they report. My data analysis is similar, but not identical, which leads me to believe they might have accidentally failed to include one or more participant surveys into their analysis or otherwise erred when calculating their values. For "ought" and "can," they report $r(80) = .18, p = .09$; I calculate $r = .189, p = .089$.²⁷⁴ For "ought" and "blame," they report $r(80) = .413, p < .0001$; I calculate $r = .401, p = .000$.²⁷⁵ And for "can" and "blame," they report $r(80) = -.07, p$

= .51; I calculate $r = -.054$, $p = .632$.²⁷⁶ Clearly the numbers are close, but not identical. There is little I draw from this discrepancy, but I thought it worth noting since one typically wants to be able to arrive at the same statistical conclusions when using supposedly identical data. I checked, double-checked, had the data double-checked in two programs and still got my results consistently, which aren't their results.

OK, but that's neither here nor there for my argumentative purposes. Up next, I look at something very important to my own study as a kind of re-creation of the original work by Chituc et al. Question by question, I took my data from the two conditions in which Brown is unable to make the meeting but has a moral obligation to do so and compared it to the statistical analysis from Chituc et al. In total, I largely replicated their findings despite my conditions containing 20 and 20 participants who took the surveys in person while thinking aloud as compared to their conditions which contained 34 and 48 participants who took the surveys online without being asked to think aloud during completion—this indicates that thinking aloud doesn't significantly interfere with how people think about these thought experiments. First, I compared their data and mine on each of three questions for each of two conditions, Brown at high and low fault. Some of these data comparisons were straightforward, others were a bit more complicated for reasons I explain below. In each case, though, since I'm hoping that my findings are similar to those of Chituc et al., I am hoping for *non-significant* differences.

High Fault, Can, Chituc et al. & My Data: I ran an independent samples t-test, equal variances assumed to get $t(66) = -.678$, $p = .500$. With that, there is no significant difference between my results and those of OIC experiment 3, which is what I should want.

Low Fault, Can, Chituc et al. & My Data: Again, I ran an independent samples t-test, equal variances assumed to get the following: $t(52) = .294, p = .770$. And again, there is no significant difference.

High Fault, Blame, Chituc et al. & My Data: For this one, a t-test didn't seem the appropriate choice. As is clear from Figure 6, my participants answered almost entirely with "5," and no one responded below 3. And while the data from OIC experiment 3 has some response lower than 3 (or 30 on their original scale), they also were highly pressed at the ceiling toward 5 (or 50 on the original scale). A more appropriate test, then, would be a chi-square test based on treating everything greater than 4.5 as one category and everything 4.5 and below as another. For the data from Chituc et al., 34 out of 48 cases were scores higher than 4.5; for my data 15 out of 20 were higher than 4.5. The results were $\chi^2(1, N = 68) = .122, p = .727$. Again, there is not a significant difference between my subjects' responses on this question and subjects' responses in OIC experiment 3.

Low Fault, Blame, Chituc et al. & My Data: I ran an independent samples t-test, equal variances assumed to get the following: $t(52) = -1.071, p = .289$. And again, there is no significant difference.

High Fault, Ought, Chituc et al. & My Data: I ran an independent samples t-test, equal variances assumed to get the following: $t(66) = -1.909, p = .061$. There is no significant difference, though it is close to a p -value less than .05. These aren't the most similar results.

Low Fault, Ought, Chituc et al. & My Data: I ran an independent samples t-test, equal variances not assumed, given that the Levene's test revealed there to be a significant

difference in variance between my dataset and Chituc et al.'s, to get the following: $t(28.191) = -1.703, p = .100$. Again, there is no significant difference.

Next, I analyzed my data between the high and low fault conditions to see if I could replicate the core finding from Chituc et al. across their three studies “that judgments of *ought*, contrary to broad philosophical assumption, do not imply judgments of *can*”²⁷⁷ and that “instead, judgments of *ought* are affected by judgments of *blame*.”²⁷⁸ In this case, I was hoping for no significant difference between conditions on the “can” question because Brown couldn’t make the meeting in either vignette, but I was hoping for a significant difference on “blame” and “ought.”

High Fault v. Low Fault, Can: I ran an independent samples t-test, equal variances assumed to get the following: $t(38) = .711, p = .481$. Here there is no significant difference, which is both desired and expected given that both vignettes make it clear that Brown is unable to make his meeting.

High Fault v. Low Fault, Blame: As with my data comparison to Chituc et al., the t-test doesn’t seem to be appropriate for “blame” question on the high fault condition given that it produced data at the high end of my 10-point scale. 15 out of 20 subjects in the high fault condition selected “5” and just 2 out of 20 did so in the low fault version. So, instead of a t-test, I calculated a Fisher exact value of .0001, which is a significance of $p < .01$.²⁷⁹ That is, the responses to the “blame” question are significantly different on the two conditions.

High Fault v. Low Fault, Ought: I ran an independent samples t-test, equal variances assumed to get the following: $t(38) = 3.218, p = .003$. And here the responses are significantly different between the two conditions.

And lastly, I wanted to look at the correlations between my responses on each question, just as Chituc et al. did. For “can” and “blame,” the correlation was $r = .031$, $p = .850$, which isn’t a significant correlation. Chituc et al., as mentioned above, reported no significance in correlating their data on these two questions either. For “can” and “ought,” the correlation was $r = -.063$, $p = .700$. Again, this isn’t significant, and it aligns with the analysis from Chituc et al. which also had no significance from these responses. And lastly, for “ought” and “blame,” there was a moderate correlation (by social science standards²⁸⁰) with $r = .482$, $p = .002$ which is significant and aligns with Chituc et al.’s finding of significance on these responses.

The upshot of all of this is that I largely replicated the findings from OIC experiment 3 and the core finding of the OIC study altogether, which was that people would judge someone to be obligated (“ought”) to carry out an action he can’t do if they judge that person to be at fault for his inability to carry out the action. Or, put more formally, Chituc et al.’s study found support for H2 over H1, and so did my quantitative data (which homed in on just two of the conditions from OIC experiment 3):

H1. Participants will deny that an agent ought to do something that the agent can’t do, regardless of whether the agent is to blame for the inability....

H2. Participants will judge that an agent ought to do something that the agent can’t do when the agent is to blame for the inability.²⁸¹

And while my focus in this dissertation isn’t on the quantitative data in OIC experiment 3, the fact that my data aligns with Chituc et al.’s data makes my subsequent qualitative analysis more compelling. I have shown that the think aloud method didn’t largely interfere with the kinds of responses participants made on the survey questions. And since my data and Chituc et al.’s data

are similar, then my qualitative data, though conducted on different subjects, speaks to their quantitative findings as well as to mine.

To end this chapter, here is a convenient index of the Appendices A – N that the reader can find at the conclusion of this dissertation.

- A. Warm-up Survey 1, “Lying”²⁸²
- B. Warm-up Survey 2, “Giving”
- C. Main survey 1, “Moral Obligation High Fault Unable”²⁸³
- D. Main survey 2, “Moral Obligation Low Fault Unable”²⁸⁴
- E. Follow-up Survey²⁸⁵
- F. Recruitment Flyer
- G. Body of Recruitment Email
- H. Quantitative Data from My Study
- I. Researcher Script Version 1
- J. Researcher Script Version 2
- K. Representative Qualitative Interview Questions Version 1
- L. Representative Qualitative Interview Questions Version 2
- M. On Transcription for My Study
- N. Additional Information of My Data Including Trivial Coding and Analysis Choices

CHAPTER SIX

A First Analysis of Qualitative Data:

Quantitative X-Phi Surveys Fail to Capture Unambiguous, Relevant Judgments

I. Full Transcripts Don't Corroborate the Ideal Story in the OIC Study

In this chapter, I am working to defend a claim specific to the OIC study: my qualitative data from think aloud and interview sessions reveal that the quantitative survey data from OIC experiment 3 point to misguided conclusions because they are grounded on quantitative methods that fail to capture unambiguous, relevant judgments. In short, the quantitative data seem to treat people's judgments on OIC, a complex moral principle, to be more straightforward than they truly are. Later I will argue that this claim indicates something about the use of quantitative surveys in x-phi more generally: quantitative surveys have limited use in x-phi if used in isolation from other research methodologies such as the think aloud method and the interview.

I want to first highlight two full transcripts that I argue demonstrate the methodological shortcomings of quantitative surveys. But before doing so, it is necessary to present the framework that Chituc et al. create in analyzing their quantitative data so I can comment on how the think aloud and interview sessions problematize their quantitative data in light of this framework. Across their three experiments in the OIC study, Chituc et al. are largely focused on testing which of two philosophical positions regarding OIC is the case. The first, in a nutshell, is that ought really does imply can in a *conceptual* or *analytic fashion*; or, more simply, "ought" entails "can" (see chapter one terminological definition of "OIC" for an extensive treatment of this). For our

current purposes, what is important is the relationship of entailment such that it becomes false that an agent ought to do something if she cannot do it, not whether this entailment is specifically one of analyticity, conceptuality, or otherwise (on the possibility that one holds there to be meaningful distinctions between analyticity and conceptuality in other contexts). The only specification on the relationship between “ought” and “can” that matters for our current purposes is that if an agent ought to do something then it is true that she can, and that if she cannot do something then it is false that she ought—this is what entailment captures (i.e., “ought” entails “can”). The specific way Chituc et al. explain the understanding of OIC that is on the table is by reference to the relationship between “bachelor” and “male”: “‘ought’ is supposed to imply ‘can’ by virtue of the concepts expressed by the words ‘ought’ and ‘can,’ just as ‘bachelor’ implies ‘male’ by virtue of the concepts expressed by the words ‘bachelor’ and ‘male.’”²⁸⁶ (Though, if one were to reject *this* comparison to “bachelor” and “male” on the basis, say, that it assumes OIC is the case based on some theory of concepts when one holds that OIC is true regardless of the conceptual relationship between “ought” and “can,” it wouldn’t matter for our current purposes as long as one understood the relationship to be one in which it is true that *agents who ought do can do* and false that *agents who can’t do ought do*.) The second view is skeptical of this claim. This skeptical view, in Chituc et al.’s work and elsewhere, is traced to work from Walter Sinnott-Armstrong.²⁸⁷ In a foundational work on OIC, Sinnott-Armstrong argues “ought” doesn’t entail or presuppose “can,” but rather (sometimes) conversationally implies it (see chapter one defining of “OIC” for more details on this).²⁸⁸ For example, ‘You ought to reboot the computer,’ a form of advising, conversationally implies that agent “can” reboot the computer, but the form of blaming “‘Where are you? You ought to be here by now!’” does not imply “can.”²⁸⁹ From this debate, Chituc et al. frame their

experiments around two empirical hypotheses, the first maintaining the analyticity of OIC and the second skeptically challenging it:

H1. Participants will deny that an agent ought to do something that the agent can't do, regardless of whether the agent is to blame for the inability....

H2. Participants will judge that an agent ought to do something that the agent can't do when the agent is to blame for the inability.²⁹⁰

I think it useful to keep their framing in mind as I analyze my qualitative data while not being fully bound by it so that I can draw other insights that relate to the concepts of “ought” and “blame” and “can” that might not fit squarely into their project.

For my present purposes, I'll note that this framing is largely unobjectionable, though I raised some concerns about it in chapter three. It seems straightforwardly true that if “ought” implies “can,” then judgments connecting these concepts in just this way shouldn't be affected by judgments of blame—or at least they shouldn't be significantly affected by them. This is the essence of H1, and therefore I think it is a valuable framing device on the relevant data: if blame judgments significantly influence judgments that should, if OIC is the case, uphold OIC, then it seems that there is a problem in asserting that OIC comports with our judgments. Of course, one could reasonably take issue with the move that more or less puts the philosophical debate surrounding OIC to an empirical test on laypeople's judgments; but for the moment, I'm not so worried about this move. (This is in part the project of Kissinger-Knox, Aragon, and Mizrahi, who take issue with the way Chituc et al.—and the same researchers in a different publication as Henne et al.—poise their empirical research as refuting OIC.²⁹¹) Likewise, H2 is straightforwardly unobjectionable in that, if shown to be the case, it really does seem to put pressure on the notion that OIC falls out of our judgments of the concepts “ought” and “can.”

Given this framing, I also want reiterate Chituc et al.'s core claim at the conclusion of their three experiments: that people's "ought" judgments were influenced by "blame" judgments and didn't simply track with "ability" judgments as OIC would predict such that OIC, as analytic or conceptual entailment, is challenged.²⁹² And to put this in the framing based on H1 and H2, this amounts to the conclusion that H2 was better supported by the data. The upshot of this is simple (at least if we assume, as I currently will, that their hypotheses H1 and H2 do help philosophers determine the truth of the principle that "ought implies can"): Chituc et al. have successfully provided empirical evidence that challenges the principle that "ought implies can" only if their participants' survey responses reflect genuine, stable judgments of "ought," "can," and "blame"; if they don't then it is difficult to see how their data discriminate between H1 and H2, or the truth or falsity of the principle that "ought implies can." And with that in mind, I want to jump into my qualitative data which I will reference to argue that there is good reason to doubt that Chituc et al.'s quantitative data reflect genuine, stable judgments of the concepts "ought" and "blame" and "can" as would need to be the case for their conclusions to be supported by their presented empirical evidence.

The first full transcript is from participant 30, who completed the high fault version. She selected that she completely agrees that Brown can't make it (score of 5), that she completely agrees that Brown is to blame (score of 5), and that she agrees that Brown ought to make it (score of 4). So how would this one participant's quantitative data, were it examined in isolation, be interpreted via Chituc et al.'s framing? In the most straightforward sense, this participant would be reporting a judgment (or multiple judgments) that Brown can't make it, that he is to blame, and that he ought to make it. Putting these items together, we might say that this looks like an affirmation that "ought" does not imply "can" since participant 30 ascribes inability and obligation

(in terms of “ought”) to Brown. Beyond this, though, we might notice that the high “blame” ascription to Brown might account for the high level of “ought” ascription, which is precisely what Chituc et al. baked into H2: that blameworthiness would swing ascriptions of “ought” to be higher when an agent is more blameworthy. So, viewed on their own and on their face, survey responses like those from participant 30 are taken to be directly supportive of H2 over H1 since H1 would predict that, despite the high “blame” ascription, participant 30 should have negated that Brown *ought* to make the meeting given that she affirmed that he can’t do so. But, as we will see, participant 30’s think aloud session and interview make it difficult to confidently argue that her survey responses reflect judgments of “ought” in relation to “can” as would be relevant to supporting H1 and H2.

As this is the first inserted transcript from my study, I want to note that Appendix M gives a detailed account of how I developed my transcript. At this juncture, there are two important features of my transcripts I want to emphasize (which are detailed further in Appendix M): 1) Because my transcripts contain dots (...) as a means of showing the connections and fluidity of a speaker’s utterances between words and phrases and sentences, I don’t use ellipses as a means to omit sections of text from my transcripts; therefore, I will only report unedited excerpts in any given quotation from the transcript. And 2) The reader should read the transcripts as if there was a blanket “*sic*” applied wholesale to them; I am confident in my accuracy in transcribing, and what might read as a clerical transcribing error is instead, as far as I can say, an atypical utterance from a given speaker.

* * *

Participant 30 Transcript

Think Aloud Session

Participant: Brown is excited about a new movie that is playing at the cinema across town. He hasn't had a chance to see it, but the latest showing is at 6 o'clock that evening. Brown's friends, Adams, ask Brown to see the movie with him, and Brown promises to meet Adams there. It takes Brown 15 minutes to drive to the cinema, park, purchase a ticket, and enter the movie. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time.

As Brown gets ready to leave at 5:45, he decides he really doesn't want to see the movie after all. He passes the time [laughs while speaking] for 5 minutes, so that he will be unable to make it to the cinema on time. Because Brown decided to wait, Brown can't meet his friend Adams at the movie by 6. OK...that's a weird way of thinking...umm... at 5:50, Brown Brown can't make it to the theater by 6. No he can't 'cause that's 10 minutes and it takes him 15...so...yea I would completely agree with that.

Brown is to blame for not making it to the theater by 6. Yes 'cause he consciously made the decision not to get there on time.

Ummm...Brown ought to make it to the theater by 6...is that saying he shou..he should have made it? [I say "However you would want to interpret it."] OK...umm...[pause]...yea I guess he should have...yea...OK.

Interview Session

Me: OK I'd like to start with question 3 since you had asked me [participant interrupts with "Mm-hmm"] if it meant one thing as opposed to another what were you thinking when you read that question?

Participant: Ummm...well I guess when I hear like someone ought to do something I think that they sh..they should do that or like it would be the right decision to do that [I interrupt with "Mm-hmm"]...umm....and so I kind of interpreted it as like Brown should try to make it to the theater by 6 [I interrupt "Mm-hmm"]...umm..so I was thinking like yea he should because he told his friend he would be there and he already knew about this rule that the theater had...umm...so that's kind of how I interpreted it yea.

Me: Was there another way that gave you pause what was what caused you in other words to...

Participant: Like question...

Me [interrupting]: Question it yea.

Participant: I think just the word ought [laughs slightly] 'cause I don't hear it [laughs slightly while speaking] very often so I think that was kind of what threw me off a little bit and like made me question the way I was interpreting it yea.

Me: If you could re-write that statement...ahhh..without the word ought how would you do so Brown...

Participant: Ummm...Brown should make it to the theater by 6 is the word I was thinking yea.

Me: And had it been worded Brown should make it to the theater by 6 what would your response be.

Participant: It would be yea he should but [laughs while speaking] he he didn't because he didn't decide he didn't want to at the last minute so I would say yea he should have.

Me: One more follow up on that [participant interrupts with what sounds like “Mm-hmm” but is somewhat unintelligible] related to question 3 do you agree or disagree with the following statement...at 5:55 Brown ought to make it to the theater by 6.

Participant: Mmmm...I would disagree with that then ‘cause I I I would think that’s that’s kind of meaning that he’s already left or like he’s he’s about to leave and he should make it there in 5 minutes versus like it says it’ll take 15 so I feel like ah..that I would that wording I would interpret slightly differently than how this is phrased ‘cause this is phrased a little bit more generally of like sh..should he like [laughs while speaking] morally get there at 6 or should he like physically like the amount of distance get there by 6 so [unintelligible].

Me [interrupting and talking over participant’s last utterance which made that utterance unintelligible]: So just to clarify the question 3 you were [participant interrupts with “Mm-hmm”] thinking a more general morally [participant interrupts and talks over my next words with “Yea”] should get there [it is also possible that rather than “get there” I say “make it there” or something similar] by 6...if I add at 5:55 [participant says “Mm-hmm”] you said physically...

Participant: Yea that would change the meaning to more like his actual like physical distance and how long it would take to get there.

Me: Is there a connection between those two meaning if he physically can’t get there [participant says “Mmm”] does that change or maintain what he ought to do in a moral sense [participant says “Mmm”] in your thinking.

Participant: Yea I feel like that does change it adding the the time and factoring that in changes it at least for me in how I think about the question yea.

Me: Would you mind elaborating on what the time does to the moral side [participant interrupts with “Yea”] of it...as best as [participant starts to speak but is cut off by my words; participant says “Eh” before being cut off] possible.

Participant: Yea I I guess the time and however far away he is is a phy...like a physical constraint that can’t really be changed but if if we’re just saying he ought to make it to the theater by 6 n...and not saying like if he leaves at this time or if he’s this far away then it doesn’t really have that constraint on it so then I think more generally versus if I was given like a specific time or distance then I would think I guess more like just practically if it’s possible...yea.

Me: Excellent. Do you agree or disagree with the following statement? At 5:40, Brown is to blame for not making it to the theater by 6.

Participant: Mmmm...[pause]...I..I guess I’m thinking that it would depend on when he decided he didn’t [laughs slightly while speaking] want to go...ummm...like if he decided at like before that time that he didn’t want to go then yes if he decided at 5:45 then I would say at 5:40 he isn’t to blame for not making it ‘cause he was still technically coulda been on time.

Me: So let me see if I’m accurately summarizing this [participant says “OK”] it’s more about when he decides I’m not gonna go I’m [participant interrupts with “Mm-hmm”] gonna wait around and less about what time it is [participant interrupts with “Yea”] when it happens [participant says “Yea”] so the blame is that conscious decision [participant says “Mm-hmm”] I’m not gonna go I’m gonna hang around [participant says “Yea”] OK great...ummm...what was your first reaction to the scenario before you got to the questions?

Participant: Umm...I think my first reaction...ummm..I kind of went slowly through the first part ‘cause the there’s uh just like a lot of information about the movie and his friend...ummm...so I was kinda jus’ tryin’ to process that part so that was my initial reaction...ummm...and then I guess reading that he decided he didn’t want to see the movie and was just gonna kind of sit around so like it wouldn’t be possible for him to see it versus like letting his friend know...umm...I guess that’s what I was thinking about was just like the situation could have been simpler er he could of at least let the other person know that’s kind of what I was thinking.

Me: When you reading the scenario were you picturing a certain time period or were you picturing that they had cell phones or that they [participant interrupts with “Hmmm”] should have used them or...

Participant: Yea I was definitely picturing like a right now type of thing yea.

Me: And if you knew for certain that they had cell phones and could communicate [participant says “Mm-hmm”] would that change any of your responses?

Participant: No I I still would stand by these statements and my reaction.

Me: Are there things a person ought to do that he or she can’t do in your thinking?

Participant: Mmmm...that’s a tough question...maybe [laughs slightly] I feel like my initial response is no...but I feel like that also contrad...might contradict what I said to this question...umm...I mean if you genuinely can’t do something I wouldn’t say that person ought to do that like if there’s a a real reason or barrier like physical or emotional or something going on then I would say no if they can’t do it just ‘cause they don’t want to do it that might I might feel different about it yea.

Me: So going with the idea that if they physically can’t do it [participant says “Mm-hmm”] you said they...something like they ought not or they it’s..is not clear that they ought [participant interrupts with “Mm-hmm”] to do it...what about Brown’s situation where he purposely puts himself in a situation [participant says “Mm-hmm”] let’s let’s make it a little bit more preventative let’s say he intentionally throws his car keys [participant interrupts with “Mm-hmm”] into the gutter [participant laughs slightly while saying “OK”] so he knows well now I can’t go [participant says “Yea”] how would you view that situation what would you think of his actions?

Participant: Ummm..I would think he he put himself in that situation it’s not circumstantial or something else happened to him that he couldn’t go so I would still feel like he he ought to have gone to the movies or should have gone ‘cause he committed to going...ummm...and that’s kind of unfair to back out at the last minute...umm..yea what wuh what was the [laughs slightly while speaking] [says something unintelligible] question [I interrupt and talk over participant’s previous words making one of the utterances unintelligible with “No no that was great”; my words here also step on participant’s next word] OK.

Me: Ummm...so one last follow up on this [participant says “Mm-hmm”] you used the some kind of reference to the past he ought to have gone [participant interrupts with “Mm-hmm”] would you in that moment the keys are now in the gutter [participant says “Mm-hmm”] would you agree or disagree Brown ought to make it to the theater by 6?

Participant: [Sighs] I would disagree at that point yea.

Me: Once it’s it already happened and he’s [participant interrupts with “Yea”] in a position where he can’t do that [participant says “Yea”]...great. Do you think the word ought can have different meanings please explain.

Participant: I think it can which is something I didn’t really think about [laughs slightly while speaking] until right now with with my own kind of hesitation in reading it and interpreting it...umm..so I guess that I interpret it to like this person should do this...ummm...I feel like it could also be interpreted as like this will happen or this person will do this...ummm...like it could be interpreted Brown will make it to the theater by 6....umm..yea so I think I think it can have different meanings.

Me: Excellent. To close, do you have any questions about the scenario or the 3 questions on the survey that you would like me to clarify?

Participant: Hmmm...I don’t think so number 3 was the only one that I had the question about.²⁹³

* * *

To reiterate, participant 30 selected 5, 5, 4 on the three survey questions regarding the high fault situation, which translates to the position that Brown can't make it, is to blame for not making it, and ought to make. Let me start by zooming in on some moments that highlight a success in the Chituc et al. survey so it doesn't appear that my standards for identifying genuine judgments are unreasonably high. When it comes to questions 1 and 2, I would be confident in claiming that I captured genuine judgments from participant 30 regarding "can" and "blame," respectively. During the think aloud session, when participant 30 comes to question 1, she states, "at 5:50, Brown Brown can't make it to the theater by 6. No he can't 'cause that's 10 minutes and it takes him 15...so...yea I would completely agree with that."²⁹⁴ She understands what is being asked of her based on the vignette, she states her reason for her response as the reason we would expect someone to have for agreeing that Brown can't make it, and she marks the survey accordingly. This is just the kind of thinking aloud that I would hope participants express when answering that Brown can't make it; this is a genuine, relevant "can't" judgment as far as I can tell. Again, with blame, we see similar clarity from participant 30 in the think aloud session: "Brown is to blame for not making it to the theater by 6. Yes 'cause he consciously made the decision not to get there on time."²⁹⁵ Participant 30 reads the question, provides a reason for ascribing blame (i.e., that Brown made a conscious decision to be late), that reason is just the kind we'd expect someone to ascribe if blame is appropriate, and she marks the survey accordingly. This is sufficient evidence for thinking participant 30 really did have a genuine judgment of "blame" in just the way that would be relevant for understanding OIC.

Again, in the interview, this blame judgment is further clarified when I ask participant 30 if she agrees that Brown is to blame at 5:40 (which is 5 minutes prior to his initiating waiting too long to make the meeting): "Mmmm...[pause]...I..I guess I'm thinking that it would depend on

when he decided he didn't [laughs slightly while speaking] want to go....ummm...like if he decided at like before that time that he didn't want to go then yes if he decided at 5:45 then I would say at 5:40 he isn't to blame for not making it 'cause he was still technically coulda been on time."²⁹⁶ In this elaboration, participant 30 ties Brown's blameworthiness directly to the moment of the conscious decision to do something that would make him late. If before 5:40 Brown had consciously decided he didn't want to go, then he would be blameworthy at 5:40; but Brown isn't blameworthy at 5:40, according to participant 30, if he decided at 5:45 since Brown would still be able to make the meeting if he left at 5:40. Again, this strikes me as sufficient evidence to conclude that participant 30 ascribed blame in a manner appropriate for researchers like Chituc et al. to take her quantitative datum relating to blame as reflective of a genuine judgment of blame as is relevant to philosophizing about OIC. The think aloud method and interview have revealed the ideal story of data collection, with respect to these questions of "can" and "blame," to be corroborated.

However, the ideal story breaks down when we get to participant 30's response to the third question, that of "ought." Again, she selected 4, which viewed in isolation speaks to her endorsement that Brown "ought" to make the meeting. In the think aloud session, she encounters the question and states the following (do note that I clarify that she interprets the question however she wants): "Ummm...Brown ought to make it to the theater by 6...is that saying he shou..he should have made it? [I say "However you would want to interpret it.]" OK...umm...[pause]...yea I guess he should have...yea...OK."²⁹⁷ Notice that she *switches* the tense of the question when she answers it, which gives rise to the subtle way in which the ideal story breaks down. Specifically, she is asked as to whether Brown "ought to" do something—that is, a present tense ascription of obligation. But in her thinking, participant 30 grapples with whether Brown "should have," which ascribes a different obligation to Brown: rather than saying he should do something,

she wonders if in the past he *should have* done something. Ultimately, she selects 4, but even as she does so, she affirms that “I guess he should have.”²⁹⁸ That is, we see a clear temporal mismatch that has important philosophical implications, philosophical implications that must be carefully articulated for their significance to be clear.

To say that “ought implies can” is to affirm that whenever someone *ought* to do something, he or she *can* do that thing; thus there is a temporal component to OIC that aligns the time of the obligation with the time of the ability.²⁹⁹ Streumer succinctly makes this point in his reply to the experimental work of Chituc et al. (in the OIC study and its appearance in the other publication as Henne et al.) noting that OIC supporters understand OIC to convey “it can be true at t_1 that a person ought to perform an action at t_2 only if this person can at t_1 perform this action at t_2 .”³⁰⁰ In line with this, I argue if Brown ought to have done something in the past, it would only violate OIC if at the time he ought to have done that thing he also lacked the ability to do that thing. Similarly, Streumer notes—with reference to experiment 1 of the OIC study, though with equal application to the presently discussed experiment 3—that the understanding of OIC captured in time-indexing (above) “allows” that we make claims about the past about what an agent ought’ve and oughtn’t’ve done *even* once agent is unable to fulfil an obligation.³⁰¹ Or in the language of this dissertation, claims of ought’ve and oughtn’t’ve (or should’ve and shouldn’t’ve) are consistent with OIC just because they don’t pit “ought” against “can” at any time. OIC, then, maps onto any temporal reference—i.e., a reference to the past, to the present, or to the future—just fine so long as “ought” and “can” reference the same time. To see this, let’s briefly examine how OIC pertains to future actions. There is nothing in the following statement that violates OIC: ‘You ought to take out the trash tomorrow.’ This is perfectly in alignment with OIC because, presumably, the obligation applies to an action that the person being spoken to *can* perform. To violate OIC in the future, as

in the past or present, we need to pit obligation and capability against each other in the same timeframe: ‘When you can’t take out the trash because you are anesthetized during your surgery, you ought to take out the trash.’ Here, the obligation to take out the trash in just the moment when the person being spoken to will be anesthetized violates OIC if this statement is taken to be true because OIC is the principle that says obligation implies ability.

When participant 30 is asked if Brown “ought to” do something, her subsequent quantitative selection would only be a reflection of a judgment relevant to OIC if she was understanding “ought” in the way that it would be pit against “can” in the way that it would violate OIC. But this appears not to be the case because she articulates that Brown *should have* done something, which is potentially a different timeframe than would be expressed by the claim that Brown *should* do something. Where does this potential time divergence come from? I argue that it is in Chituc et al.’s questions themselves. Question 1 asks about Brown at 5:50, effectively asking if Brown has an ability at a very specific time, but question 3 makes no specific temporal reference, leaving open the possibility that some readers *hold fast* to 5:50 but others interpret the question regarding Brown’s obligation as less temporally bound. This wouldn’t be a problem if Chituc et al. were arguing about *how* people interpret Brown’s obligation as such, but they aren’t. Instead, they are leveraging claims against OIC, a principle which only pertains to instances when there is temporal consistency between the obligation and the ability. In other words, their third question on the surveys that I employed in my study lend themselves to an opaque notion of “ought” but their conclusions treat it as non-opaque such that they understand quantitative endorsements of “ought” to indicate a judgment of “ought” which is relevant to affirming or undermining OIC, a principle that is only affirmed or undermined by judgments based on temporally pitting “ought” against “can.” Of course, were think aloud sessions and interview to

reveal that participants largely understand the vague temporal language of “ought” to be *that which applies at 5:50*, then we would have confirmed the ideal story Chituc et al. rely on in making their unique claims about OIC. But, as I show below, and as I have hinted at in my preliminary look at participant 30’s think aloud transcript, no such reality is the case. Before looking at that, though, I want to point out that Leben, in a study I’ll discuss in more detail next chapter, notes something similar with relation to experiments 1 and 2 in the OIC study, arguing that Chituc et al.’s work insufficiently accounts for the problem of “temporal underspecification,” which in essence is a lack of clarity as to whether participants are being asked about agent’s obligation before or after they become unable to fulfil it.³⁰² As I’ve just noted with experiment 3, this underspecification problematically leads researchers to misinterpret what participants might be expressing in their quantitative responses.

Participant 30, in saying that Brown *should have* done something appears to be articulating the judgment that Brown, *prior to being incapable of meeting his obligation or generally speaking, ought to have* done those things that would have had him be on time for the meeting. It doesn’t appear to be an endorsement that Brown at 5:50, the moment he was incapable of making the obligation, *ought to* meet his friend. In the think aloud session alone, I think we have enough to speculate that this is what appears to be participant 30’s judgment. And this speculation alone is enough to undermine any confident claim-making on behalf of data from people like participant 30. But with the interview in addition, we can actually draw out the judgment in more detail—I heard participant 30’s temporal discrepancy, and I seized on it in the interview. I argue the ideal story doesn’t hold given that participant 30 elaborates in so many ways to show that she views present obligation and past obligation very differently, thus deflating the idea that her quantitative data reflect something like an anti-OIC judgment. This is most clear in the large section below,

worth reading in full to get a sense of how it unfolds. Notice that participant 30 appears to treat “ought” differently when a concrete time is brought into the question about Brown’s obligation:

Me: One more follow up on that [participant interrupts with what sounds like “Mm-hmm” but is somewhat unintelligible] related to question 3 do you agree or disagree with the following statement...at 5:55 Brown ought to make it to the theater by 6.

Participant: Mmmm...I would disagree with that then ‘cause I I I would think that’s that’s kind of meaning that he’s already left or like he’s he’s about to leave and he should make it there in 5 minutes versus like it says it’ll take 15 so I feel like ah..that I would that wording I would interpret slightly differently than how this is phrased ‘cause this is phrased a little bit more generally of like sh..should he like [laughs while speaking] morally get there at 6 or should he like physically like the amount of distance get there by 6 so [unintelligible].

Me [interrupting and talking over participant’s last utterance which made that utterance unintelligible]: So just to clarify the question 3 you were [participant interrupts with “Mm-hmm”] thinking a more general morally [participant interrupts and talks over my next words with “Yea”] should get there [it is also possible that rather than “get there” I say “make it there” or something similar] by 6...if I add at 5:55 [participant says “Mm-hmm”] you said physically...

Participant: Yea that would change the meaning to more like his actual like physical distance and how long it would take to get there.

Me: Is there a connection between those two meaning if he physically can’t get there [participant says “Mmm”] does that change or maintain what he ought to do in a moral sense [participant says “Mmm”] in your thinking.

Participant: Yea I feel like that does change it adding the the time and factoring that in changes it at least for me in how I think about the question yea.

Me: Would you mind elaborating on what the time does to the moral side [participant interrupts with “Yea”] of it...as best as [participant starts to speak but is cut off by my words; participant says “Eh” before being cut off] possible.

Participant: Yea I I guess the time and however far away he is is a phy...like a physical constraint that can’t really be changed but if if we’re just saying he ought to make it to the theater by 6 n...and not saying like if he leaves at this time or if he’s this far away then it doesn’t really have that constraint on it so then I think more generally versus if I was given like a specific time or distance then I would think I guess more like just practically if it’s possible...yea.³⁰³

While the nuances are always difficult to draw out in a transcript alone, I feel comfortable detecting that participant 30 seems to treat the original question about if Brown ought to make the meeting to be a general question as to whether it is morally right for Brown to make the meeting. And, to this, participant 30 seems to agree. And this is good in showing that participant 30 finds that Brown has a moral obligation to uphold his promise, but it says nothing about OIC which has a unique temporal component to it. For her responses to negate OIC, it must be that she wants to affirm at 5:50 or 5:55 that Brown can’t make it and that he ought to. But when 5:55 is brought in—5:55 being a time when Brown *clearly* can’t make it—participant 30 no longer affirms that

Brown ought to make it, noting that the physical constraint of the situation changes the meaning of “ought.” Whereas the think aloud session presented only an intimation of a perspective on ought other than the one needing to be captured for H1 and H2 to be weighed, the interview granted me the opportunity, through targeted follow up questions, to delve deeper into this intimation and reveal that participant 30 really doesn’t seem to judge that when Brown is in a state of non-ability that he is also held under the obligation. The main idea I want this to evidence is that cases like participant 30’s show that Chituc et al.’s quantitative data, when viewed in isolation, lead to potentially misleading conclusions. Here we have someone who would be taken as undermining OIC doing nothing of the sort, and potentially reifying it by being careful *not* to affirm Brown’s obligation at 5:55 given material constraints on his behavior.

Before moving on to the next full transcript, I want to note that we see both the think aloud session and the interview at work here, both providing us with useful information about participant 30’s judgments. The think aloud session gives us not only her thinking as she encounters a given question, but it also gives me a moment to flag for elaboration in the interview so I can follow up with relevant questions. And then, with the interview, I am able to dive into more detail regarding a participant’s thinking which I wasn’t able to do as they are thinking aloud in the first place. Importantly, the think aloud method and interview proved to be useful in those cases when the ideal story was corroborated *and* in those when it wasn’t. And again, this is the main point I hope to be making: the MPEP is such that we lack epistemic justification for taking quantitative data alone to be reflective of the philosophical judgments we treat them as reflecting, but with qualitative methods we can both provide this justification (when we were in fact justified in doing so) and reveal that we lacked it (when we were in fact lacking it). I just can’t help to see how comforting it is to see participants, like 30, clearly articulate their reasons for selecting a response

when those reasons are the ones we hope they draw from—such as with “can” and “blame” judgments. And, I just can’t see how someone could stay comfortable making any real claims of substance based on survey responses that appear to have been selected for reasons different from those that would be relevant to understanding the philosophical question at the heart of a given survey—such as with “ought” judgments. I think what we have here, and in the other examples I provide, is not a judgment of “ought” relevant to undermining OIC, but instead we have something superficially similar.

To see similar corroboration and breaking down of the ideal story, I want to look now at the second full transcript, which is from participant 14, who completed the high fault version. She selected that she completely agrees that Brown can’t make it (score of 5), that she completely agrees that Brown is to blame (score of 5), and that she neither agrees nor disagrees that Brown ought to make it (score of 0).

* * *

Participant 14 Transcript

Think Aloud Session

Participant: Brown is excited about a new movie that’s playing at the cinema across town...umm the name Brown strikes me as kind of surprising if it’s a first name. He hasn’t had a chance to see it, but the latest showing is at 6 o’clock that evening...which also seems strange to me that usually there are later showings than that...umm...Brown’s friend, Adams...so maybe I’m wondering [unintelligible] they’re both last names...asks Brown to see the movie with him, and Brown promises to meet at..Adams there. I’m re-reading that just to make sure I got it so...Adams is inviting and Brown says yes I’ll meet you there...OK...it takes Brown 15 minutes to drive to the cinema, park, purchase a ticket, and enter the movie. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time...wondering about previews ‘cause it seems like most movies have a significant chunk of previews before...umm...as Brown gets ready to leave at 5:45, he decides he really doesn’t want to s...wait get’s ready to leave...so I’m assuming that means like wherever he was before he went to the movie...umm..he doesn’t want to see the movie after all. He passes the time...oh OK...so I’m anticipating that he’s he’s gonna..hhh...blame the movie theater for the fact that he doesn’t meet his friend...umm he pa passes the time for 5 minutes, so that he’ll be unable to make it to the cinema on time. Because Brown decided to wait, Brown can’t mate...meet his friend Adams at the movie by 6...K...at 5:50...do you agree or disagree with the following statements...at 5:50 Brown can’t make it to the theater by 6...so I’m looking back to just to check it takes him 15 minutes to drive there and park and do all that...umm...so as long as...make it to the theater means like into the movie theater f.. to see the movie then I think that I completely agree so I’m selecting plus 5.

Umm...Brown is to blame for not making it to the theater by 6. Yes I completely agree because he made an active choice to wait when he he could have left and would have made it in time...umm...Brown ought to make it to the theater by 6. So I feel like you can’t really contemplate that unless because of the the way the verb tense is you you

can't contemplate unless it's like if up to 5:45 if it's after that an..then it's already too late but...umm...I...hmm...I'm th...I'm thinking like it's a crappy thing to do to your friend because your friend has now spent the money and is seeing this movie and a lot of people don't like to see movies alone...umm...I love seeing movies alone so there's part of me that's a little biased...umm...it's not like you left your friend to do something horrible by himself but umm...I mean yea you you probab...that was kind of crappy...uhh..I'm going...I'm between 0 and plus 1...umm..I think I'm gonna s...the verb tense is is throwing me but if like if I'm imagining that I'm his head at the moment...in in his head at the moment when he's making this decision...I'm actually going to say 0...I just don't think the stakes are that high for this...umm...OK.

Interview Session

Me: To begin, I'd like to pick up with question 3 [participant says "Mm-hmm"] and I was wondering if you could elaborate on your thinking there you s..you said something about the time at 5:45 [participant interrupts with "Mm-hmm mm-hmm" over my next words, too] you mentioned a few other things if you wouldn't mind as [participant interrupts with "Sure"] much as possible unpacking that...

Participant: Umm..so my my training is actually in in English grammar..umm..an ESL teacher by background and I work in writing now and so I noticed that it says ought to make it which is like a like a present tense thing an..i..I noticed that it didn't say should have made it...umm...so that threw me a little bit...umm...and I and I felt like I can't I can only think of the answer to this question if I'm like in that moment where he..im...with him in that moment trying to decide if I'm if I'm as has him going to try to make it or not...umm..does that answer your question? [I say "Mm-hmm"] yea..OK.

Me: So to follow up you said in that moment what moment are you thinking of?

Participant: I think the moment like right around 5:45 where he's saying I have to leave right now if I'm going to make it...and he's deciding if he's going to take the steps to to make that happen or not..so presumably like the minute or two before 5:45.

Me: And had the question been worded in reference to the past [participant says "Yea"] Brown ought to *have* made it...would that change your response?

Participant: No I don't think it would...uh..well...let me think...Brown ought to have made it to the theater by 6...no I think I would still say neither ag..agree nor disagree...umm...I I the the word ought or should for me like I said before...is like it's it's a very strong thing and this to me as a kind of moral or ethical decision d..like just doesn't seem like the stakes are that high and the th there's not gonna be much harm caused if he doesn't go...so I think it's not a very nice thing to do to your friends but it isn't like a huge moral issue for me.

Me: And one last follow up on that question, the selection of neither agree nor disagree [participant says "Mm-hmm"] if y..is as much as possible e..even if that is reiterating a bit what what made you choose that as opposed to agreement or disagreement?

Participant: Partly because if I agree..uh..I don't see a reason not to agree completely..like just to agree a tiny bit...to me the situation just doesn't have the sort of things you'd.. the sort of boxes you'd need to check off for the stakes of it to matter enough that I would like have a strong feeling either way...umm...if it said something like it would be considerate for him to have [I interrupt with "Hmm"] made it I I would I would definitely agree with that but ought...umm..maybe I have like a higher standard of of what needs to be met in order for that to be like a mandated uh uh like a really should have happened that way.

Me: And had it said..had the stakes been higher that might have also changed your answer?

Participant: [sighs]...Like if his friend was waiting for a ride to the airport or something and he did this it would have totally changed my answer [I interrupt with "OK" as participant says next words] because that would create some harm to the friend like they would have to spend more money or they would be late for a flight or something but in

this case basically the friend sees the movie that they were going to see anyway and just sees it alone which I don't consider to be a great hardship.

Me: OK, perfect...thank you [participant says "Yea"]. What was your first reaction to the scenario you read about with Adams and Brown?

Participant: That...umm...I'm really introverted and I often find myself when I've made plans with someone to do something like going to a movie or and eh eh...any kinda kind of social thing...I often find myself last minute feeling like why did I think that [laughs while speaking] I was [unintelligible]...so I can kind of maybe that influences my sense that he didn't do something so bad...umm...I was also..uhh...I was surprised by this scenario I didn't expect it to be something like this I think maybe because the other ones felt a little more like this kind of lying feels pretty maybe morally questionable to me and then the the the donating thing it feels like there's there's some level of harm and so this one it seemed a little more benign.

Me: Great [participant laughs]. What did question 1 on the survey ask?

Participant: It seemed to me like it was a factual question and it's just asking you to carefully read the information and based on what I've been told he won't..there's there's not enough time in 10 minutes to to do something that takes 15 minutes...so I don't think that there's a need for the like the variations of agreeing or disagreeing...umm..I guess it's possible that he could have like he could have speeded sped [laughs]...umm.. or like skipped some stop signs or something and gotten there sooner but it it just seems fairly [unintelligible] it's just like a yes or no question to me.

Me: What did question 2 on the survey ask?

Participant: Has to do with culpability...umm..and whether he's responsible for the choice and the result of that choice which I think he clearly is.

Me: And what was your reasoning for agreeing on number 2?

Participant: Because he made a conscious choice and the choice made something no longer a possibility and no one else was able to make that choice in this moment and so I think he's solely responsible for it.

Me: Great. Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: [laughs] So I'm...I don't think I c..do I have to agree or disagree..I don't think I can answer that because that also like on a grammatical level for me implies that he's on the way and the question is like is he gonna make it and like yea you know he he should be able to make it that that's how I hear that question...[I say "OK"]..can you read it again?

Me: At 5:55, Brown ought to make it to the theater by 6.

Participant: I don't think that we have enough information to answer that..or I have enough infor..[unintelligible]...the way I'm conceiving of it the question and the situation I don't think I have enough information to answer that.

Me: If we knew that at 5:55 Brown is still at home...

Participant: Then no..he can't. [I say "He can't..."] He can't make it by 6 because if it's true that it takes 15 minutes then then no.

Me: OK, perfect. Are there things a person ought to do that he or she can't do? Please explain.

Participant: Hmm..[pause]..well..it raises the question of whether you can whether it's legitimate to say that someone ought to do something that they're incapable of...um...I think people might perceive that someone ought to do something that they're incapable of doing...umm...if it's physically impossible then I don't think it's reasonable to

say that they ought to do it...if i...I'm thinking of like if you have the emotional resources or the energy to do something and you really might be struggling and so maybe there is something that you should do but you make a choice that you're...you don't feel in the moment able to do it and that that feels like those things can kind of be in tension but if it's literally impossible like if someone is paralyzed and they see a person who's drowning or whatever it you know like we could say a human should try to save a drowning person but that person's not capable and so I yea I'm not sure whether that person ought to then...it's just like a generic person ought to do this but if you if you don't have the ability to do it it's not your fault [laughs].

Me: Perfect. The word ought is an important word in the survey. Please describe how you understood the word ought when it was used in question 3.

Participant: Umm...I initially understood it as like it would have been the right thing for him to do like the the the good, morally appropriate thing for him to do...uhh...that that's what it was implying and asking me if I agreed with that umm..except for the verb tense thing which was throwing me...umm...it it occurred to me later that there's another way to read it which is like guessing whether something is likely to happen just sort of logistically like what I said about maybe he's on the way and he's almost there and I can say like oh he ought to make it by 6 but I think actually usually ought has more of a sort of moral sense whereas should can have that kind of probability or likelihood sense slight like a slight difference in meaning although they're pretty similar.

Me: Just so I want to make sure I'm understanding these the the way you were using it when you were saying probability [participant says "Yea"] is it something like based on where we are we should be there by 6...[participant says "Yes, exactly"] OK and the other sense you said was something with moral..[participant interrupts "It's the right thing to do..yea it's moral...morally right"]...so would you agree that there are different meanings of the word ought..[participant interrupts with "Sure, yea" and talks over my next words] or do you disagree? [participant says "I I agree"]

The scenario does not mention cell phones, and it does not mention whether or not Brown tries to contact Adams [participant interrupts and speaks under my words with "Hm-mm"] if you knew that both Brown and Adams had cell phones to communicate on [participant says "Mm-hmm"], would you view the scenario differently?

Participant: So I was actually assuming that they had cell phones to communicate on so I would say like if I knew that they didn't then I would view it differently.

Me: And how so?

Participant: Because if you're just gonna leave the person there and like they don't know why you didn't come or whatever that's pretty that that's worse [laughs]..umm...but if you have the ability to let them know that you're not able to come then I then then tha that was my assumption and so I answered it based on that...umm...does that...did I just answer your question? [laughs]

Me [interrupting the laughter from participant]: Yes, that's perfect. To close, do you have any questions about the scenario or the three questions on the survey that you would like me to clarify?

Participant: [pause] I'm wondering if the the ought to was intentionally in the the present tense or or if it if it was meant to say like do you think he should have tried to or...

Me: So the specific reason why it was phrased that way [participant says "Uh-huh"] I don't know but I know that it is designed to see how someone interprets [participant says "OK"] so there is possibly multiple ways that someone could interpret it [participant interrupts with "Got it"] or maybe they view it as just it very clearly means this [participant interrupts with "Yea yea"] and so it's really open to [participant says "Uh-huh"] the reader [participant says "OK"] to see how they [participant interprets with "OK, got it" and talks over my next words] interpret it.

Participant: OK. I think that's it.³⁰⁴

* * *

I want to note again we see corroboration on the questions of “can” and “blame”—we can be comfortable saying this participant reported relevant judgments in both cases. In the think aloud session, participant affirms that Brown can make it: “at 5:50...do you agree or disagree with the following statements...at 5:50 Brown can’t make it to the theater by 6...so I’m looking back to just to check it takes him 15 minutes to drive there and park and do all that...umm...so as long as...make it to the theater means like into the movie theater f.. to see the movie then I think that I completely agree so I’m selecting plus 5.”³⁰⁵ And, this is further iterated in the interview:

Me: Great [participant laughs]. What did question 1 on the survey ask?

Participant: It seemed to me like it was a factual question and it’s just asking you to carefully read the information and based on what I’ve been told he won’t..there’s there’s not enough time in 10 minutes to to do something that takes 15 minutes...so I don’t think that there’s a need for the like the variations of agreeing or disagreeing...umm..I guess it’s possible that he could have like he could have speeded sped [laughs]...umm.. or like skipped some stop signs or something and gotten there sooner but it it just seems fairly [unintelligible] it’s just like a yes or no question to me.³⁰⁶

In both we see a consistent response—that Brown can’t make it—and consistent reasoning—there’s not enough time. I take this to be a solid corroboration of the quantitative results. For blame, we see similar clarity. First, in the think aloud session: “Umm...Brown is to blame for not making it to the theater by 6. Yes I completely agree because he made an active choice to wait when he he could have left and would have made it in time.”³⁰⁷ And again iterated in the interview:

Me: What did question 2 on the survey ask?

Participant: Has to do with culpability...umm..and whether he’s responsible for the choice and the result of that choice which I think he clearly is.

Me: And what was your reasoning for agreeing on number 2?

Participant: Because he made a conscious choice and the choice made something no longer a possibility and no one else was able to make that choice in this moment and so I think he’s solely responsible for it.³⁰⁸

With clear articulation of just those reasons that we’d hope, participant 14 by all outward appearances expresses a genuine, relevant judgment of blame. In both these cases of ability and blame, I say we’ve corroborated the ideal story Chituc et al. are after.

But with “ought,” there is no such corroboration. We first get a sense of participant 14’s “ought” judgment not being the kind relevant to OIC in the think aloud session:

Brown ought to make it to the theater by 6. So I feel like you can’t really contemplate that unless because of the the way the verb tense is you you can’t contemplate unless it’s like if up to 5:45 if it’s after that an..then it’s already too late but...umm...I....hmm...I’m th...I’m thinking like it’s a crappy thing to do to your friend because your friend has now spent the money and is seeing this movie and a lot of people don’t like to see movies alone...umm...I love seeing movies alone so there’s part of me that’s a little biased....umm...it’s not like you left your friend to do something horrible by himself but umm...I mean yea you you probab...that was kind of crappy...uhh..I’m going...I’m between 0 and plus 1..umm..I think I’m gonna s...the verb tense is is throwing me but if like if I’m imagining that I’m his head at the moment...in in his head at the moment when he’s making this decision...I’m actually going to say 0...I just don’t think the stakes are that high for this...umm...OK. ³⁰⁹

At the start, participant 14 states that the verb tense is such that it is difficult to answer, only to articulate something that looks like an endorsement of OIC: “unless it’s like if up to 5:45 if it’s after that an..then it’s already too late but.”³¹⁰ The answer to the question about “ought” appears to be different if participant 14 interprets it such that she is being asked about Brown *up until* 5:45 than if she interprets it as *after* for the very reason that he can’t make it after 5:45. But then, participant 14 goes on to highlight a different issue altogether: that little is at stake affects her judgment of ought. Clearly, the selection of “0” in the end captures neither the first line of reasoning, which appears to *endorse* OIC, nor the second, which brings in an entirely different issue, namely whether ought applies more or less to more or less consequential events. Instead, a

neutral score is read in the data analysis as less of an endorsement to H2 than a score of 1 or higher, but more of an endorsement for H2 than H1 since a failure to disagree with “ought” isn’t supportive of OIC (or H1) which would require just that kind of disagreement given that Brown can’t make it. In the interview, though, we can see an expanded dialogue on this issue and we learn that participant 14 might very well endorse OIC despite what her quantitative score indicates:

Me: Great. Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: [laughs] So I’m...I don’t think I c..do I have to agree or disagree..I don’t think I can answer that because that also like on a grammatical level for me implies that he’s on the way and the question is like is he gonna make it and like yea you know he he should be able to make it that that’s how I hear that question...[I say “OK”]..can you read it again?

Me: At 5:55, Brown ought to make it to the theater by 6.

Participant: I don’t think that we have enough information to answer that..or I have enough infor..[unintelligible]...the way I’m conceiving of it the question and the situation I don’t think I have enough information to answer that.

Me: If we knew that at 5:55 Brown is still at home...

Participant: Then no..he can’t. [I say “He can’t...”] He can’t make it by 6 because if it’s true that it takes 15 minutes then then no.

Me: OK, perfect. Are there things a person ought to do that he or she can’t do? Please explain.

Participant: Hmm..[pause]..well..it raises the question of whether you can whether it’s legitimate to say that someone ought to do something that they’re incapable of...um...I think people might perceive that someone ought to do something that they’re incapable of doing...umm...if it’s physically impossible then I don’t think it’s reasonable to say that they ought to do it...if i...I’m thinking of like if you have the emotional resources or the energy to do something an n you really might be struggling and so maybe there is something that you should do but you make a choice that you’re...you don’t feel in the moment able to do it and that that feels like those things can kind of be in tension but if it’s literally impossible like if someone is paralyzed and they see a person who’s drowning or whatever it you know like we could say a human should try to save a drowning person but that person’s not capable and so I yea I’m not sure whether that person ought to then...it’s just like a generic person ought to do this but if you if you don’t have the ability to do it it’s not your fault [laughs].³¹¹

What seems clear is that participant 14 talks about what a person ought to do in very different ways when there is a concrete temporal component linking it to their ability to do so—e.g., when asked if Brown ought to make it at 5:55 when he is still at home, she asserts, “Then no..he can’t”³¹²—as compared to when there *isn’t* a concrete temporal component linking it to their ability—e.g., “we could say a human should try to save a drowning person but that person’s not capable and so I yea

I'm not sure whether that person ought to then...it's just like a generic person ought to do this but if you if you don't have the ability to do it it's not your fault.”³¹³ Unlike her straightforward expressions about “can” and “blame,” participant 14 offers a complicated view on “ought” that doesn't clearly get captured by her 0 score as it sits in conjunction with her selections of 5 and 5 on the first two survey questions. And, more interestingly, the view expressed seems to fit more squarely with an endorsement of OIC, not its negation.

Again, with participant 14, we see the think aloud method and interview providing both corroboration and debunking with respect to the ideal story of data collection. The think aloud method provides some hints of judgments more nuanced than those available in the quantitative scales, and the interview allows sufficient space and time for those nuances to be drawn out in more detail. This further iterates the value of these methods, as well as highlights the misleading conclusions we draw if we assume quantitative survey data is valid because it is face valid.

II. Further Analysis of the Ideal Story from Selected Excerpts

I'm less interested in showing corroboration on questions 1 and 2 about “can” and “blame,” respectively, because they aren't as interesting to the project of understanding OIC, which is really about understanding obligation more so than ability. Here are a handful of excerpts that display the corroboration or breaking down of the ideal story with relation to “ought” judgments beyond the full transcripts of participants 30 and 14.

The first excerpt is from participant 4. She selected 5 on all three questions for the high fault version, which should indicate that she is supportive of H2 as she affirms both Brown's inability to make the meeting and his obligation to do so. In her interview, it becomes clear that

her initial response of 5 in agreement that Brown ought to make the meeting reflects a more general sense of “ought” such that it is not pit against Brown’s inability at 5:50 to make the meeting:

Me: So what do you think about the following statement? Do you agree or disagree with this following statement. At 5:55, Brown ought to make it to the theater by 6.

Participant: Uhhmmm....I think [unintelligible] would really depend on where he was. [laughs]

Me: According to the passage, where...?

Participant: No. I don’t agree then. ‘Cause 5:55 he would need 15 minutes to get there. I’m still assuming he’s at home, in my head [I say “Mm-hmm”] yea.

Me: So if we’re assuming he’s at home at 5:55 the statement Brown ought to make it to the theater by 6...just to reiterate...what was your response?

Participant: No. Completely disagree.

Me: And why..why would that change it from agree to disagree?

Participant: Huh. [pause] ‘Cause for me Brown ought to make it to the theater by 6 it’s more of a responsibility thing like yea he really shoulda done that, but if you’re putting a time on it like if he left at 5:55 he ought to make it by 6 that’s a time constraint so that doesn’t make sense ‘cause he can’t *physically* do that..you know..does that make sense?³¹⁴

Notice that the material constraint on Brown’s behavior at 5:55 is what causes her to change her response from completely agree that he ought make it to completely disagree that he ought do so. This shows that participant 4 is judging “ought” and “can” in a way that is *supportive* of OIC rather than negative of it. And yet, with quantitative data alone, Chituc et al.’s framework would have me viewing her data as just the opposite.

The second excerpt is from participant 5, who completed the low fault version and selected responses 5 (Brown can’t), 0 (Brown is to blame), -4 (Brown ought). With such responses, we should expect that she endorses the principle of OIC or at least doesn’t report anything to negate it. And, in large part, I think this is corroborated:

Me: What did question 3 on the survey ask?

Participant: Brown ought to make it to theater by 6.

Me: And again, what did you choose?

Participant: I chose negative 4 umm..almost completely disagree because...do you want me to say why? [I say “Mm-hmm”] OK, yea. Because umm..he’s not gonna make it to the theater by 6 umm...what regardless of whether or not he should be there to meet his friend...umm..I think the question is heavily weighted because his friend was planning on meeting him there but if he was going by himself then it wouldn’t be as big of a deal about his car breaking down.

Me: Do you agree or disagree with the following statement: At 5:55, Brown ought to make it to the theater by 6.

Participant: [pause] No...I disagree..um..because there would definitely not be enough time.

Me: Excellent. [pause] Would you mind elaborating on you said there definitely wouldn’t be enough time and so how does that connect with the statement that Brown ought to make it to the theater by 6?

Participant: Umm...so that it at...so instead of Brou...[unintelligible] ought to make it by 6...it’s it’s 5:55 and he ought to make it by 6. [I say “Mm-hmm”] Umm..[sigh]..the first..the question number 3 doesn’t have a time—oh I guess I never really thought about it like that—so the first question doesn’t have a time...so Brown ought to make it to the theater by 6 I guess that could be just a general question regardless of what happened with his car or even what transportation he chose to take so like no matter what he ought to make it to the theater by 6 like that could be true..but now if we’re putting the 5:55 it sounds like all the other factors that happen to him are incorporated and he’s definitely not gonna make it to the theater by 6 or just in general if he left at 5:55 he would not make to the theater by 6 so in that case I would completely disagree that he can’t make it.³¹⁵

Participant 5, in the multiple ways she discusses obligation and ability, centers on the former being negated when the latter is negated. And while she never comes out and says “ought implies can,” participant 5 appears to be judging something like that. In this case, I would feel confident saying that she had reported a genuine, relevant interpretation of “ought” as would be relevant to philosophizing about OIC because she provides her reasoning for disagreeing with Brown ought to make it as the fact that he can’t make it; this is reiterated again when I put a time constraint on the question. And so, this looks like a corroborated judgment in support of OIC. Again, whether the ideal story is corroborated or challenged by the think aloud method and the interview, it shows the value of such methodologies. Given that some quantitative data are corroborated and others not, it casts doubt on how much we can make of the quantitative data as a whole.

The last excerpt is from participant 13 who was given the low fault version and selected 5 (Brown can’t), -5 (Brown is to blame), and -5 (Brown ought). These quantitative responses seem to indicate that participant 13 either endorses OIC or doesn’t negate it outright. And, I argue that we do have that corroboration in very important moments from the interview:

Me: And what did question 3 on the survey ask?

Participant: Yea that was kinda odd I thought...that he ought to make it...so then I had to back and check the details [laughs]...like could he make it in time...ummm..I also think maybe at that point even if I was off on the on the math or the time that that's why I said this guy's having a bad day because it seems like at that point..I I wouldn't blame him for calling his friend and just saying you know this...uh...I gotta get my car towed or let's just re-schedule...

Me: What did you choose...

Participant [interrupting]: So I didn't think he ought to have to do that it sounded like what responsibility did he have...yea.

Me: Would you mind elaborating on you said it the the last question was odd or it was y...

Participant [interrupting]: Well it made me think that he..umm..the first one was just almost like a math question you know...can he make it in time...so that's why I double-check my the figures since and thought well did I miss something like is this a trick question [laughs]...[I say "Mm"]...[laughter continues]...and then then the second and third were more like does he have some personal responsibility and the third one implied that maybe he could but that he was not going to...and that would you know stand up his friend and be mean to him who's waiting at the theater for him...so that made it seem more like he was actually hurting a person...the ought is what implied that to me.

Me: When you said he could jus' in a reference to number 3 you said it made it sound like he could, what do you mean by that?

Participant: Umm..[whispers unintelligibly]...oh... um um...he's to blame for not making it to....well I didn't th...I d..well I don't remember when I said could... umm..that umm...yea well the second one I didn't think it was his fault because of what happened to him...ummm...but the last one did imply to me that maybe he could 'cause it was saying ought that he ought to make that well maybe he could...you know maybe he should and then you know again I was checking to make sure that...umm...I got the figures and time right...

Me [speaking over participant who says "Yea"]: And just to clarify maybe he could maybe he should do what exactly?

Participant: Ummm...meet his friend at you know go to the theater...I mean I see this all about 'cause you know the only person he has to be responsible for other than his friend is himself and if he doesn't go he doesn't go for whatever happens...but he's got another person involved so I was thinking that you need to think about this person Adams who is...umm...gonna meet him there and is he gonna stand up his friend you know so I was always thinking of how would this affect Adams.

Me: If instead of his car breaking down [participant says "Mm-hmm"] he decided to just wait around until 5:50 [participant says "Mm-hmm"] so that he didn't have enough time to make it would that change your response to number 3? So the car breaking down you said it wasn't his f..[participant interrupts "Oh"] but [participant says "If he just...yea, it would, it would"]...wh..so what would your [participant interrupts "Ummm"] response be for number 3...

Participant [interrupting]: Yea if he just decided at...what he waited around until 5:50 [I say "Or even 5:55 until his window is..."]...yea I mean..umm..you know yea he should have then..ah ah ah...I to me 'cause you're..you've made plans with a friend and yea I feel like you have a responsibility to that person if you don't let them know or or something happens but yea then I think he should have..yea [I interrupt and say "OK"] yea...absolutely...yea.

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: At 5:55? Umm... I feel like I'm missing something here [laughs]...[I say "No just however you are interpreting the questions and the scenario" of which the last words are talked over by participant's next words]...I know but I'm I'm I'm thinking...uh...dh..ch.ch.ch.ch.ch...[whispers unintelligible words]...no I mean I don't think he could...I don't think he could cause it says here that his car...umm...he can't fix it himself in time and it's too late to make it by bike so...I mean I don't think so unless I'm missing..how could he get there on time.

Me: Would it matter if at 5:55 he was there because of the car breaking down or his own choices...would that make a difference...at 5:55 time [participant interrupts with “Mm-hmm”] is running out [participant talks over my words with “Mm-hmm, mm-hmm”] ought... ought he be there by 6?

Participant: Mm-hmm...ummm..yea I mean I think that would make a difference yea I think that would make..again it comes back down to if it was his fault or he just was being selfish and ‘cause that’s what it seems like to me that you’re....it’s just like well somebody’s waiting for me and...you know I’ll just because if you get there and it’s too late then you not only made unless your friend already goes in which we don’t know...ummm..then umm...you know you’ve inconvenienced somebody else.

Me: Perfect. Are there things a person ought to do that he or she can’t do? Please explain.

Participant: [whispers something unintelligible]..oh ought and can’t [laughs]...must be somebody’s philosophy [both laugh]...OK...ummm...if you ought to do that you can’t do..oh I don’t know how you do that...if...um...I I don’t know how you do that I I I see where if you don’t want to do something but you ought to do it but if you can’t do it like if I can’t contribute monthly like Josephine did then..ummm...I don’t know that I..how could I ought to do it I mean I maybe I should do something else that I can do [I say “Mm-hmm”] like you know you can volunteer at a soup kitchen or something if you’ve want to help but I I don’t see how you can do that...they seem like they work against each other.³¹⁶

First, note that participant 13 reads the third question about “ought” as *implying* something about Brown’s ability: “and the third one implied that maybe he could but that he was not going to...and that would you know stand up his friend and be mean to him who’s waiting at the theater for him...so that made it seem more like he was actually hurting a person...the ought is what implied that to me.”³¹⁷ That is, the very question about “ought” *implied* that Brown could make the meeting. This closely connects with OIC. In addition, when she does affirm that Brown ought to do something, she reverts to the past participle form which, as previously discussed, removes the temporal link needed for this agreement of “ought” to be a negation of OIC: “Yea if he just decided at...what he waited around until 5:50 [I say “Or even 5:55 until his window is...”]...yea I mean..ummm..you know yea he should have then.”³¹⁸ This interviews reads as if participant 13 can’t bring herself to say something that violates OIC, even to the point where she assumes she must be missing something because I keep asking about something that seems obviously true, namely that when Brown *can’t* make it we can’t say he *ought* to:

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: At 5:55? Umm... I feel like I’m missing something here [laughs]....[I say “No just however you are interpreting the questions and the scenario” of which the last words are talked over by participant’s next words]...I know but I’m I’m I’m thinking....uh...dh..ch.ch.ch.ch.ch...[whispers unintelligible words]...no I mean I don’t think

he could...I don't think he could cause it says here that his car...umm...he can't fix it himself in time and it's too late to make it by bike so...I mean I don't think so unless I'm missing..how could he get there on time.³¹⁹

Participant 13 links Brown's ability and obligation so closely, it seems, that when she is asked about whether he ought to make it at 5:55, she reiterates that he can't.

Participant 13's apparent upholding of OIC is clear enough that it comes through even for a specific moment in this exchange that I want to address, where for the briefest of seconds it looks like she says one thing that could have momentum toward a negation of OIC. Specifically, I start a question with "would it matter if at 5:55 he was there because of the car breaking down or his own choices" in what is most obviously an effort to give her the opportunity to ascribe "ought" just in the case that Brown is more blameworthy, and then I end the long-winded, and not greatly worded, question with "ought he be there by 6?"³²⁰ And to this blurb of a question, participant 13 begins her response with "Mm-hmm...ummm..yea I mean I think that would make a difference" such that she might want to say that Brown oughtn't when he is not to blame but ought when he is, but it is clear in the remainder of her response that her focus is on *the content* of Brown's actions and their *consequences*, not on ascribing "ought" to him specifically at 5:55.³²¹ In fact, her response attends almost entirely to some kind of difference in Brown's behavior, not between whether or not "ought" at 5:55 really applies in low versus high blame scenarios. She certainly doesn't take the opportunity to emphasize something like 'Well yes, if he made the choice then even though he can't make it at 5:55 he ought to.' When her response is explored in the greater context of her dialogue, I think it actually speaks to her *support* of OIC, for that when she is offered an opportunity to lay obligation on the blameworthy Brown, she bypasses it as if the true question I was getting at had to do with something else—i.e., whether consciously choosing to opt out of plans is different than being victimized by car troubles—possibly because the question about whether Brown "ought" at 5:55 is, as explored in her previous utterances, almost confusing. But

just to send it home, we can note that her reply to the very next question about OIC has her affirming it once more: “if you ought to do that you can’t do..oh I don’t know how you do that...if...um...I I don’t know how you do that I I I see where if you don’t want to do something but you ought to do it but if you can’t do it like if I can’t contribute monthly like Josephine did then..umm...I don’t know that I..how could I ought to do it.”³²²

Again, the qualitative methods help us corroborate the ideal story with this participant, which is supportive of OIC rather than in opposition to it. And we’ve also now seen that some cases where support for OIC shows up in the quantitative data, the opposite seems true in the transcript data. I’m not concerned with generalizing about laypeople’s judgments toward OIC, but rather I hope this shows that the quantitative data is sometimes accurate and sometimes not, so it really can’t be relied on in isolation.

III. A Reluctant Summary of My Qualitative Data, or “Did Anyone Clearly Negate OIC?”

Without occupying too many pixels, I want to emphasize that I am reluctant to give a general summary of my qualitative data. Then I will give that summary.

First, as is clear in this project, I’m skeptical of the air of objectivity surrounding quantitative or summative claims. We see, some of us, a high percentage of participants support this or that response and it seems to convey that a high number of people have the judgment supposedly captured by that response which subsequently puts pressure on the philosophical claim supported by that judgment. To read “87% of participants” intuitively, to me, reads as firm, grounded, and decently objective. But, as I hope to have made clear, this can be misleading. Second, the numbers game is especially problematic in χ^2 where (an ultimately misguided) criticism is that philosophical truth isn’t determined by vote. In the face of such an objection,

Knobe and Nichols make it clear that x-phi surveys have philosophical value indirectly³²³—experimental philosophers aren't asking you to reject your theories because a majority of laypeople do. Nonetheless, I am wary of placing much stake in quantitative values as they can potentially morph into implicit justifications for a kind of activity that makes x-phi susceptible to such criticism, namely refusing to shoulder a burden of proof because one's preferred theory has majority support. Third, quantification of participant support for this or that response naturally draws our attention to the majority or plurality view rather than some other view which maintains sufficient support so as to warrant philosophical attention. If 10% or more of participants express some view, this tells me there's something philosophically interesting going on with the view that oughtn't be dismissed because it doesn't make up the majority position. So, I am cautious because there is risk of ignoring the views less ascribed to once quantitative data is placed at the center of an investigation. Fourth, my goal in this project is to use experimental methods to better understand how people think about ethical issues, namely those bound up in OIC, and subsequently inform the expert debates on said issues on behalf of my findings and musings. And, I feel I can most clearly and helpfully do this without becoming ensnared in the narrowly defined debate as to whether OIC is or isn't the case, or whether people do or don't find it intuitive, because I think the experimental data consistently points to how complex people's judgments are—something I'll continue to unpack in the next chapter—such that it would be premature to reduce my work to a defense or challenge to OIC. And lastly, and relatedly, the very richness of people's ethical thinking on OIC that gives me pause in drawing firm conclusions about how laypeople generally think about it, in combination with the messy nature of qualitative data, adds up to the reality that it is quite difficult to give a precise sense of how many participant responses supported or negated OIC.

Still, I want to give some sense of what I found in the think aloud sessions and interviews as it pertains to the intuitiveness of OIC. In reading over the transcripts in full, earmarking relevant utterances and exchanges and then later re-reading those earmarked notes in context, I found that most things people said were consistent with or supportive of “ought implies can” either before or after the necessary temporal concreteness was made salient, and few participants said anything that seemed to clearly negate OIC in the face of temporal concreteness being clarified. Often participants superficially seemed to negate OIC in their responses before the timing was clarified, but once clarified they would utter something compatible with or in support of OIC. Also, participants would often uphold consistency with OIC by switching the language from “ought to” (and the like) to “ought to have” (and the like). For example, consider this exchange with participant 22:

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: Uhh...does that mean he's leaving his house by 5:45?

Me: Let's imagine the scenario follows like this at 5:45 he's decides I'm just gonna [participant says “Mm”] hang around and then end up not going so it's 5:55 he's still at the house he hasn't left...

Participant: Oh..well...in that case he can't make it there so he still ought to have been at the theater by 6 but umm...yea I guess at that point it's not possible same with the like...yea..it's just not possible.³²⁴

The superficial interpretation might have this as a negation of OIC, but upon closer inspection it's clear that her switch to language of the past, as to what ought to have transpired, preserves OIC. But in other cases, some I've already highlighted and others I will highlight next chapter, participant claims appear to support OIC more straightforwardly.

To put it succinctly, the participants whose transcripts I highlight in this chapter and the next aren't atypical in their responses in apparent support of OIC; I chose them but I could've chosen others to include. I ended up using these participants, though, because their responses were especially clear and interesting in their understandings of ideas related to “ought implies can.” My

overall impression of the data, then, is that most participants upheld or spoke in line with “ought implies can” as an intuitive principle, and they certainly *didn't* consistently say things that clearly violate it like, ‘I know he can't do it, but he still ought to; the ought doesn't go away because in that moment he can't do it any longer.’ And they did have *plenty* of opportunities to offer up such negations. I think it so important to not only appreciate the many utterances supportive of OIC but also to appreciate what *we should've seen* were OIC to be counter-intuitive in Brown's situation(s). If OIC weren't intuitive, we'd expect more participants to consistently, effortlessly, uncompromisingly offer statements such as the following: 1) ‘At 5:55 Brown can't make it, but his obligation is still there. He ought to but can't’; 2) ‘The obligation doesn't go away just because your car breaks down. Even at 6 he ought to be there despite the fact he can't’; 3) ‘Brown chose to make it impossible after 5:45 but he still ought to make good on his promise; he can't make it because of his choices but he ought to do, even once it's too late.’ But in my reading the transcripts and re-reading my notes, I don't see that—and though it's of course possible I missed something like this in one or more transcripts, it is highly unlikely. To make this clearer, let's imagine I was asking participants about a principle that we all agree is false: “ought” implies “hat” (OIH), which is the principle that agents who ought to do something are necessarily wearing a hat, and if they aren't wearing a hat then it is false that they ought to do that thing. Imagine that instead running into car troubles or deciding not to keep a promise, Brown is on his way to keep his promise to Adams with enough time to do so when at 5:45 he takes off his hat in the car. Then imagine I asked participants if they agree with the following statement: “At 5:55, after his hat is off, Brown ought to meet Adams at the movie by 6.” I think it's clear people would be happy to say the obligation holds regardless of whether he's wearing a hat, saying things like “I agree because his hat has nothing to do with this; he is still obligated to go” or expressions of confusion at such an

absurd question. If OIC were really that counter-intuitive, then that truth would come across as clearly (or close to it) as would the counter-intuitiveness of OIH.

That said, I do want to highlight a few instances where it appears OIC is negated *in some sense*, though it will be obvious that some such cases aren't as clear-cut as we'd like. I want to now highlight a few examples of utterances that seemed to negate OIC, few as they were. And, to be fully transparent, there were a few other instances not included that were quite fuzzy (and fuzzier than those below) but might have captured some essence of the negation of OIC.

Perhaps the best example comes from participant 7 whose transcript in some moments speaks to OIC, and in others a negation of it. Participant 7 dealt with the high blame case and responded Brown can't (4), Brown is to blame (5), and Brown ought (3). Most negative of OIC is this moment since it is after I've clarified the temporal layer of the "ought" question:

Me: So let's assume that the passage when it says it takes him 15 minutes that that is definitively how long it would take on the best case scenario [participant says "Mm-hmm"]. I think you're right that it's not clear about this is [participant says "Mm"] *absolutely* how long it takes [participant says "Right"] it..the only cutoff we get is 6 pm [participant says "Mm-hmm"]..let's assume though he really needs 15 minutes [participant say "Mm-hmm"] otherwise he will be late [participant says "Yup"]..so running with that, do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: Disagree. [I say "Would you mind elaborating on why?"] Umm..he needs 15 minutes to get there, assuming that 5:55 he hasn't started the journey yet there's no way empirically if it takes 15 minutes at the best to get there, there's no way that by leaving at 5 he's gonna get there on time, so why ... why start? On the flipside, it..that would be almost the ri..that he that he he's giving himself..ehh..at 5:55 again he has an excuse in place for..umm...for not taking..for not taking responsibility..umm..but so if if he waited until 5:55 he he's still..OK that's [laughs] that's interesting 'cause if he waited until 5:55 he ought to go I would say based on ...wh wh..based on the fact he's standing up his friend his friend's counting on him being there there was a contract in place which he is breaking, yes he still ought to go..uhh. ...I agree with that but if even if he goes nothing's gonna change, the contract's broken..he Ada...he's not he he still gets what he wants, Adams doesn't..g..swh.. get what he wants..so..[sighs]...umm..so I guess...yea he he he he oughta still keep his commitment, but I understand at 5:55 there's no way the commitment's going to be kept so why why why why do that, why why step out and do that?

Me: Great, so on the survey and and this is to [participant says "Mm-hmm"] kind of elaborate on what you just said, on the survey you had put initially with the statement Brown ought to make it to the theater by 6 you said agree [participant says "Mm-hmm"]..when we put the time at 5:55 you initially said disagree [participant says "Mm-hmm"]..as best as possible if you could elaborate why that response would change.

Participant: I think because..I I I w..wo.. in terms of..fo..for lack of a better term in terms of the the moral dilemma that I see with Brown..Brown made a commitment to go to the movies with Adams he knows he has to get there at 6 o'clock..umm..so yes in that regard he ought to make it to 6 he ought to keep his word keep his contract that way every that way both both sides get what they want..but..at 5:45 as he's getting ready to leave the contract changes and all of the sudden he doesn't want to see the movie and by going to the movie he's gonna get he's he's gonna get something

he doesn't want he's gonna have to do something he doesn't want...ummm...and then...umm..breaks it by..umm...by waiting so [sighs] so yea he he ought to make it to the theater by 6 ...um...and that's why I p..why I put agree..umm.. because like keep your commitment keep the contract..umm..you know keep your word..umm..but by but again by waiting and by you know 5:50 5:55 you know still not out the door..umm...it's just the the whole notion of whether whether it's even worth it anymore kinda changes because if he leaves he'll still he can still try but he's not gonna get there so...why try to begin with so..I I still agree he ought to make it to the theater by 6 empirically but once that changes and once there's other parameters in place like OK well 5:55 should he still go..well no the damage is done and can't be undone by him going so therefore you know that kinda changes the answer a little bit to disagree because like now that it's 5:55 well you're not gonna fix it by going therefore you know..you oughta make it but you're not going to so don't...don't try, so it it shifts a little bit more to the to the disagree at that point.³²⁵

While his first response is supportive of OIC, there are clear moments later in the responses that seem to negate it. In particular, participant 7 seems to be saying that the obligation holds despite Brown's inability, but it would be pointless for Brown to do anything about it.

The next clearest is from participant 23 who deals with the low fault condition—and provides responses Brown can't (2), Brown is to blame (-5), Brown ought (-5)—but helpfully speculates regarding a scenario very much like the high fault condition in her think aloud session:

Brown ought to make it to the theater by 6....um....I feel like if it were any other reason aside from his car from breaking down...umm...and he was just tryi...being flaky then I would feel stronger about him committing to his promise or to his umm...arrangement with Adams but ought to make it to the theater by 6 given these circumstances I'm gonna say completely disagree negative 5.³²⁶

Then, in the interview, I seek clarification:

Me: To begin I'd like to start with the third question you said because it was something out of his control you put completely disagree but you had mentioned something like that might change if other circumstances [participant says "Correct"] would you mind elaborating on that [participant interrupts with "Sure"] or or correcting if I didn't summarize that correctly.

Participant: Umm...yea..ul...well I feel like if it's really the result of extenuating circumstances external things that are preventing him from actually making it and he has the the desire to go but actual logistics can't make it...umm...I feel like I would disagree with the statement but if he were more just flaky or...ummm...he didn't manage his time well then I would be more prone to saying that he should've followed through on his commitment.

Me: So let's imagine that rather than his car breaking down Brown decided I'm gonna wait purposely uhh...until 5:55 [participant says "Mm-hmm"] knowing that I can't make it if I leave by then [participant says "Mm-hmm"] so he waits around till 5:55 so that he won't be able to make it [participant says "Yea"] w..would you agree or disagree with the statement Brown ought to make it to the theater by 6.

Participant [interrupting and talking over the last of my previous words]: I would agree...ummm...I would I would probab..ahhh...he knows that the...umm...cinema I'm assuming that he knows that the cinema has this strict policy...ummm...of not admitting anybody and he has 5 minutes to get there and there's no mode of transportation that can get h..him there in 5 minutes I would probably lean to a plus 4.³²⁷

What would have made this clearer is just if she articulated why she holds Brown still obligated here but not when he isn't culpable, and also if she hadn't seemed so focused on the timing of his travel—for a second it seems like her answer is contingent upon whether he might be able to make it. Also, there is a reading of this that looks like she wants to crank up the agreement with “ought” as if to say ‘Brown better get things in gear!’ rather than he ought do what he can't do—notice her use of “should've.”³²⁸ Still, this looks like a pretty clear violation of OIC.

The third is from participant 2³²⁹ (from whom I examine an excerpt in the next chapter), who selected that Brown can't (2), Brown is to blame (2), and Brown ought (-5) on the low fault survey—a combination that speaks to a consistency with OIC. And the overall sense from his transcript is that he upholds OIC; consider this exchange as context:

Me: What do you think about the following statement: At 5:55, Brown ought to make it to the theater by 6.

Participant: Uhh..I'm thinking there's, number 1 there's not enough information if his car's still broken down, if it is there's no way, I'm thinking there's no way. But it he... it gives him even less time to actually make it and so I'm thinking I would agree that he could not make it by 6.³³⁰

But to be charitable to the anti-OIC position, one moment reveals a possible negation from participant 2:

Me: Great. So the word ought is an important word in the survey, please describe how you understood the word ought when it was used in question 3.

Participant: Umm..yea...it..it is an interesting word and I...I see the word ought as ummm...an expectancy like they ought to know better or I sh..er...he..he ought to be there by a certain time or ummm..but I see it as a definitive term ummm...and faith in other people, in other..in other...umm..the way other people go about their business. Personally, I would only use ought in the sense of like I'm assuming that they would know something like they *ought* to know something or they *ought* to be there. So there's a surety that comes with it. So, I guess for question 3, Brown ought to make it to the theater by 6, going back, like, he.. he ought to make it there but I...personally.. I don't I don't think he cou...he could make it there by 6 just based on having less time and .and..ummm...you know..the distance to the theater.³³¹

The last bit has a ring of “ought” doesn't imply “can,” but it also looks like participant 2 is saying that yes Brown ought to keep his promise, in some ideal sense, but that participant 2 in particular

doesn't think he can. When placed in the larger context of his transcript, during which he responds with -5 that Brown ought, this moment really could go either way in support or negation of OIC.

Next, an exchange with participant 9 produces what is superficially a very clear negation of OIC, but in broader context isn't so clear:

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: Umm...initially I would say neither agree nor disagree because I don't know if he's in the vehicle I don't know where he is or what his intention is.

Me: If we assume that after deciding to not go [participant says "Mm-hmm"] and waits around for 5 minutes until the window..um..is less than 15 minutes to make it and he makes no effort to go...at 5:55, Brown ought to make it to the theater by 6 what do you think of that statement?

Participant: I would say completely disagree because..um..we don't know well we don't how long it takes to get to the movie theater because he has no intention of making it I would imagine that even if he was in his car for whatever reason or was like making movement..umm... yea there's not strong enough attitude towards making it to the movie for him to make it there make it to the theater by 6 or we don't have reason to believe so.

Me: Great. And you mentioned the word intention..with..uh..with r...re..relation to the third question. How are you uh connecting intention with that question would you mind elaborating?

Participant: Yea..so..um..I am guess I can look at intention as [sighs]..um..like attitude..um...mixed with maybe beliefs towards doing something or behaving in a particular way..um..and so ought is a..um..what it's a ..um..yea I because of the word ought..um..I am losing my word right now but..um..oh it's like right..[unintelligible]..it's right on the tip of my brain [I say "Mm-hmm"]...umm..yea..I..OK so I guess where why I'm using the word intention is..umm..ought to make it making it sound like it might be up to Brown..um...to make it to the theater by 6 and because he now does not have the desire to go to the theater..umm..which I suppose desire could have been used synonymously with intention..um...in the question..yea..that's why I used the word intention...does that answer [I say "Mm-hmm"] your question?

Me: What if the the phrase rather than using ought said at 5:50 Brown is obligated to make it to the theater by 6?

Participant: Umm..yea I would say that like that sentence would n...ahh...make a lot of sense..umm..is obligated to m...at 5:50 Brown is obligated to make it to the theater by 6...umm..Brown has let's see let me read back the first paragraph..Brown is excited about a new movie that's playing at the cinema across town he hasn't had a chance to see it, but the latest showing is at 6 that evening. Brown's friend Adams asks Brown to see the movie with him, and Brown *promises* to meet Adams there. So again this is like maybe..um.. looking at how we how I would imagine...um...or what I define a promise to be and I do see a promise as something that is held to be true and it is an oblig..or it cou..can be seen as an obligation depending on what the promise is..so yea I think if you would have..if the number 3 would have asked umm..Brown is obligated to make it to the theater..um...by or is in the car and is obligated to make it by 6..um...then I would have said completely agree. That question it wouldn't 'cause it doesn't imply whether he's going to it is only talking about..uhh..his relationship to to Adam..um..and de..and creating a promise with Adam.

Me: Great, and one more follow up on that if we added uh..a an even tighter time constraint and said at 5:59 Brown is obligated to make it to the theater by 6.

Participant: Umm...yes..I would keep the same..yea the time doesn't matter it's more about the obligation in the question.³³²

In an effort to make sense of how participant 9 uses “ought” I asked about obligation, specifically. This negates OIC in that she responds by saying that Brown is obligated in virtue of making a promise, and she supports this position even when asked about 5:59 when Brown clearly can’t make it. However, the exchange reads a bit as if she takes me to be asking simply if there is an obligation applied to Brown in the scenario, which she confirms based on the promise that was made. Once this view is snapped into frame, then the remark that “the time doesn’t matter”³³³ reads more like she is saying that the time is irrelevant in determining whether Brown formed an obligation here. It is beyond the scope of this project to deal with the language of “obligation” versus that of “ought,” but I’ll say this: even though philosophers may talk about OIC in relation to obligation, I think there is reason to think that being asked about “obligation” primes in some people different judgments than being asked about “ought” since obligations, depending on the context, can act as social rules that apply to certain relationships that function differently than moral ought’s. For example, someone can talk about being obligated to attend a work function that they can’t attend as a result of being out of town, and yet this doesn’t violate OIC: if they can’t attend (then on OIC) it is not the case they ought to attend while it may be the case that they are obligated by virtue of working for a given company. Nonetheless, this exchange may very well negate OIC.

And lastly is a brief exchange from participant 34 (whose interview I look at in the next chapter), which is only slightly in support of a negation of OIC. More than anything, there’s a combination of words that reads like an expression of “ought” and “can’t,” but I hesitate to make much of it given the strong support we find for OIC in the rest of her transcript:

Me: So for question 1 you agreed that at 5:50 Brown can’t make it to the theater by 6 [participant says “Mm”] then for question 3 you agreed that Brown ought to make it to the theater by 6, looking at these 2 responses together do you agree or disagree that Brown ought to do something that he can’t do?

Participant: Mm...well the question of ought to is dependent on what basis you're you're looking at...umm...and so if you're looking at the plan that they had he ought to follow the plan if it was agreed to and...umm...by both parties...umm...now since things have essentially he's essentially changed his agreement in the you know uh very loose contract or plan that they have...umm..what I think is a better question is that there ought to be some more communication..umm...that could therefore change the plan and create a difference...umm...ummm...list of oughts basically...umm...and so at the moment we don't have any information whether he's communicating at all and so if this is all the information that is true and present then yes he finds himself in a conundrum of what he ought to do is something he can't do and I think that's where a lot of people find this procrastinating paralysis [laughs]..umm...because they're possibly cornered between a rock and a hard place that being of guilt and capability.³³⁴

In negation of OIC we get the language of Brown being stuck in a sense. But that too is what makes this not so clearly a negation since “ought” obtaining in light of inability isn't a claim about a person's psychology being jammed up; it is a moral claim about obligations holding for unable agents. The emphasis on feeling stuck, therefore, makes this read like Brown is experiencing the unpleasant results of no longer being able to carry out some action. Still, this moment speaks to a negation of OIC.

In total, more often people spoke in terms compatible with or even supportive of OIC than in opposition to it. Even if it were just those participants in this chapter and the next saying things in support of OIC out of a total of 40 participants, that would be enough for us to reconsider the claims of Chituc et al. and others who find little intuitive support for OIC. But of course, as I've summarized, there were even more participants that seemed to endorse OIC or speak in ways that preserve it. At the same time, there is at least some negation of OIC in my transcripts, but it wasn't common or as clear-cut as I'd expect if OIC were counter-intuitive. Still, it is enough to draw our consideration. Perhaps there is some portion of a population who views OIC as not so conceptually or analytically connected; and their views too should inform our discussion of the philosophy of OIC.

IV. How this Examination Indicates a General Limitation of Quantitative Surveys

I lastly want to reiterate that both the think aloud method and interview proved useful in both corroborating the reporting of genuine, relevant judgments *and* in showing the breaking down of the ideal story of data collection. And, for both uses, these methods prove themselves to be useful to x-phi research.

While I have focused on the OIC study by Chituc et al., the think aloud sessions and interviews seem useful in relation to almost *any* x-phi project wherein judgments/intuitions are being captured via quantitative surveys. It might often be the case that these methods reveal that the ideal story was corroborated, I have the sense that it will often reveal this not to be the case just in those places where the interesting judgments appear. Is it any surprise that judgments about Brown's ability are more easily corroborated than whether or not he has a moral obligation? Both vignettes I use (high and low fault) even assert that Brown can't make it on time! But when the question turns to "ought," we oughtn't be surprised that participants offer varied views that express confusion or nuance or uncertainty; moral obligation is conceptually complex. Would OIC be an interesting debate in philosophy were there some straightforward way to view it much as we can straightforwardly assert Brown can't make it at 5:50 if his car breaks down? If one can catch the intuition that what makes x-phi worth conducting are those issues for which there is ambiguity and complexity, then hopefully one can also catch the intuition that those are the very moments where we should expect to see these qualitative methods as useful and revealing. Though I focused on OIC in particular, a key debate in moral philosophy, it seems intuitive to me that these qualitative methods would be useful even if they merely corroborated the ideal story of other studies, but that almost certainly they will do more than that in revealing where and when it breaks down.

CHAPTER SEVEN:

A Second Analysis of Qualitative Data:

The Think Aloud Method and Interview Suggest that OIC

Judgments Are Sensitive to Concreteness, Definitional

Distinctions, and More

I. My Study in Connection with the Experimental Philosophy of OIC

In the previous chapter, my goal was to show that viewing quantitative survey data from OIC experiment 3 in isolation leads to misleading conclusions about people's judgments. I wanted to show that quantitative x-phi surveys fail to consistently capture relevant, coherent philosophical judgments. From there, I made the case that while I only examined one experiment, the kinds of limitations I exposed in the survey data of OIC experiment 3 are almost certainly not unique to it and therefore x-phi studies in general would benefit from incorporating qualitative data such as those collected from the think aloud method and the interview. In total, I argued against the use of surveys in isolation in x-phi. In this current chapter, my goal is to analyze the qualitative data so as to posit some plausible accounts of the many nuanced contours of people's understandings of the relationship between blame, obligation, and ability. Specifically, I want to connect my data with the experimental philosophy of OIC, which represents a recent development in the ongoing use and debate around "ought implies can" and its significance in moral philosophy.

The larger debate surrounding "ought implies can" focuses on the relationship between moral obligation and an agent's ability to uphold it. To understand this relationship, philosophers grapple with questions of modal and deontic logic, physical and psychological ability, determinism

and free will, blame and moral responsibility, knowledge and moral responsibility, cognitive processing and moral obligation, and much more. In other words, the superficially-simple three-word principle ends up touching on some of the most far-reaching and important ideas in moral philosophy, political philosophy, legal philosophy, epistemology, metaphysics, philosophical logic, etc. Recently, the larger OIC debate has been tackled by experimental philosophers, and this has given rise to an especially interesting sub-debate that has more focus on the following: the intuitiveness of OIC, whether its intuitiveness bears on the analyticity of the moral principle, the role it plays in social cognition, and how we might come to further grapple with moral obligation in light of empirical data capturing laypeople's employment of the concepts of "ought," "blame," "can," and their corollaries. It is this small pocket of the OIC debate that I am entering into in this section, though the conclusions of this debate certainly link up with the larger issues philosophers care about regarding moral obligation.

This sub-debate of OIC, which I will refer to as the x-phi of OIC, is framed in a few different ways, but we can generally talk about it as trying to get at the following questions: 1) Do laypeople judge the relationship between "ought" and "can" to be one that coincides with the principle OIC? 2) Do people find OIC to be intuitive? 3) What are the different ways people understand and employ concepts related to moral obligation, moral responsibility, and ability? 4) Do people understand concepts of obligation and blame and ability in a variety of ways? 5) How do people understand obligation and ability in relation to different language of obligation and ability—e.g. how do they think of expressions of "ought" as a form of advising v. a form of expressing blame? 6) How do people understand concepts of "ought" and "ability" when agents are unable to carry out a moral obligation, especially in complicated situations where, for example, an agent loses her ability due to unforeseen circumstances on the one hand or as a result of

purposeful self-sabotage on the other? Importantly, these questions aren't viewed as being solely confined to moral psychology, but instead they are viewed as being important for understanding the philosophical principle of OIC itself. If OIC lacks intuitiveness, then we might reconsider the analyticity of OIC. Or, if it is intuitive, we might fully abandon any worries about its analyticity. Also, if people tend to understand obligation in a variety of ways, it might tell us that the relationship between obligation and ability is more complicated than it is supposed. In sum, the sub-debate borne out of x-phi research makes contributions to moral psychology *and* to moral philosophy; it gets to the heart of the purely philosophical debate surrounding OIC by drawing from empirical data about people's judgments/intuitions regarding "ought implies can" and related concepts.

One important link between the larger OIC debate and the sub-debate grounded in x-phi research has been the work of Walter Sinnott-Armstrong, specifically his thought experiments wherein an agent makes a promise to meet a friend but then either falls victim to external forces that prevent agent from making the meeting or agent consciously chooses to put himself in a position wherein he is unable to meet the obligation.^{335, 336} These are the same kinds of cases explored in the OIC study by Chituc et al. (of which Sinnott-Armstrong is a co-author). And, at the heart of their project is an exploration of Sinnott-Armstrong's work, an exploration as to whether he is correct in arguing "that 'ought' does not necessarily, analytically, or conceptually imply 'can.'"³³⁷ Such a viewpoint motivates the Chituc et al. study and highlights the contributions x-phi research can make to the larger OIC debate: if laypeople employ concepts related to OIC in ways that support Sinnott-Armstrong's thinking, then it seems like he is on to something missed in the debate, namely that something assumed to be intuitive and conceptually true might not be

so. But if people tend to uphold an analytic link between “ought” and “can,” then there is more philosophizing needed in order to make sense of Sinnott-Armstrong’s case.

What seems agreed upon in the x-phi of OIC debate is that empirical data of laypeople’s intuitions are relevant to the larger OIC debate. But precisely how is not universally agreed upon. For example, Henne et al. (the same researchers as Chituc et al.) write as if their data support a refutation of the analyticity of OIC³³⁸; others argue that certain data challenge just the *intuitiveness* of the principle, which should caution anyone endorsing the philosophical principle based on an appeal to intuition. Specifically, Kissinger-Knox, Aragon, and Mizrahi take Buckwalter and Turri (in their work “Inability and Obligation in Moral Judgment”) along with Chituc et al. to task over how we ought interpret empirical studies that appear to undermine OIC.³³⁹ In the end, they argue that “the correct conclusion to draw from the empirical fact that people attribute moral obligations to agents even when those agents are unable to fulfill those moral obligations is that OIC is not intuitive, not that OIC has been refuted,” which means philosophers who maintain its truth “can no longer treat it as an axiom.”³⁴⁰ I don’t think it necessary to decide on precisely which position is more accurate in order to move forward in looking at data on people’s judgments/intuitions on OIC, because all views center on a value in understanding how people employ these concepts. Better yet, I’m going to take a different approach altogether and try to contribute to the debate based on the following positions, each of which I label for convenience:

1. ***Ought Non-controversially Implies Can:*** In some uncontroversial sense, it seems obvious that moral obligations are often if not always concerned with things agents can actually do. The Ten Commandments unsurprisingly appeal to actions agents *can* do, ostensibly³⁴¹; Peter Singer asks us to consider the life you *can* save, not the one you can’t³⁴²; Aristotle endorses virtues one *can* embody³⁴³; Kantian ethics deals with

actions people *can* typically do, like making and avoiding false promises.³⁴⁴ Little (to no) moral philosophy is dedicated to holding people morally obligated to do things that are both straightforwardly moral and absurdly impossible, such as successfully willing one's body to no longer require sleep so that one can spend an additional 6 hours each day building moral communities and helping care for people in need of food, medicine, etc. In this general sense, OIC as a principle isn't going anywhere. Or, to put it more eloquently, Kurthy and Lawford-Smith, in a piece I explore later, claim that "no one thinks that 'all oughts imply can'"³⁴⁵ in response to an x-phi study that fails to adequately account for the ambiguity of "ought." That is, if this or that study successfully finds that people are happy to attribute obligation and inability to an agent, that doesn't obviously wipe away a principle which "holds about as much consensus as it is possible to get in a discipline characterized by disagreement."³⁴⁶

- 2. *The Attempted Counter-Examples In Some Sense Reify OIC:*** If you want to see if OIC is intuitive, then it seems odd that you would focus so intently on cleverly constructed situations wherein agents seem to be exploiting a kind of loop-hole to the rule—e.g., Brown intentionally waits too long so he can't meet his obligation to Adams—rather than focus on myriad other cases where the intuitiveness seems warranted. That's not to say that these loop-hole thought experiments aren't useful; quite the opposite is true. But the intuitiveness of a principle shouldn't stand or fall on such oblique cases. We'd hardly find the moral principle that "murder is wrong" to be less intuitive if the utilitarian crafted a clever thought experiment that complicates this general principle such that the lives of millions of innocent people might hang in the balance. No, but such utilitarian thought experiments, like Sinnott-Armstrong's cases

in the OIC debate, expose a deeper complexity to a generally intuitive and useful principle and therefore warrant exploration. In reality, the fact that the debate as to whether OIC is intuitive or not comes down to cases like the one in which Brown self-sabotages or some external force prevents him from making his meeting actually speaks to *the otherwise intuitive nature* of OIC.

3. ***Sinnott-Armstrong and OIC Skeptics Have Produced Something Valuable:*** While OIC generally doesn't seem mutable, we have the invitation to learn more about the *complexity* of OIC in the examples from Sinnott-Armstrong and others. That is, philosophical discussion around OIC is strengthened by skeptical positions toward OIC and clever thought experiments that make such skeptical positions more plausible (or at least not easy to explain away). Sinnott-Armstrong's cases, and those derived from them, put OIC under the microscope in really interesting ways without asking the philosopher or the experimental participant to know much about the surrounding OIC debate: the details are simple enough to follow, the situations are realistic, and the tension between obligation and ability is palpable. By exploring such cases, we stand to learn a great deal about obligation, blame, and ability in unique ways, as we contemplate how to make sense of obligation when inability arises because an agent self-sabotages or because he succumbs to external forces. I doubt we'd learn as much about obligation and ability if we focused solely on situations wherein an agent is blamelessly rendered unable to carry out an obligation.
4. ***New X-Phi Research is an Invitation to Explore:*** I argue that my study shows that, as expected, when an otherwise intuitive claim like OIC is attempted to be "broken," we get to learn interesting ways about how people understand obligation and ability and

blame. My contribution, then, is to highlight some interesting moments in my data that I think point to possible ways people understand these concepts, and to point to promising ideas for further empirical exploration. I am not interested in evaluating the truth of OIC or in proposing a unified theory of it; both are beyond the scope of this project and would require substantially more empirical research. I don't have a position on whether we are evaluating the intuitiveness of OIC or its analyticity, but there's clearly more to be explored *by way of* capturing and analyzing people's judgments and bringing those into the x-phi of OIC.

5. ***This Exploration is Best Achieved via Triangulation:*** In reading the empirical studies of OIC, I consistently see an unsubstantiated ideal story, an over-reliance on quantitative surveys even when a qualitative free response is included, a jumping to conclusions, and room for qualitative methodologies (more on this below). We have a complicated topic on our hands, and I think the back and forth between quantitative and qualitative and mixes of both will help us better understand how people understand concepts of obligation, blame, and ability.

With these founding positions in mind, let's now start to connect my study with this debate in the x-phi of OIC.

In the next section, I will outline my positive contributions to the debate within the x-phi of OIC. To do so, I will further analyze data from my study while weaving in relevant claims and findings from the surrounding literature. In the remainder of this section, though, I want to give a brief overview of this literature so it is clear how my work fits into it. It offers many great insights that I haven't the scope to discuss, so I will keep my overview mostly to the core features of the empirical experiments therein and their implications. In addition to the OIC study and the

piece by Henne et al., which presents data from the OIC study, there are 9 publications that constitute the core literature on the x-phi of OIC. I don't summarize them in chronological order, but in a strategic way designed to show a birds-eye-view of the unfolding conversation around OIC.

The first I'll outline is the previously mentioned work by Kissinger-Knox, Aragon, and Mizrahi entitled "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations'"; beyond doing a great job in highlighting the important distinction between empirically challenging OIC itself and its intuitiveness, this piece doesn't add much to my current discussion that the OIC study hasn't already brought to the table. The authors do conduct a study on OIC, but it is entirely quantitative in nature, and the administered surveys contain enough room for varied participant interpretation that I find it succumbing to the MPEP quite clearly.³⁴⁷ They did find that participants attributed moral obligation to agents even when they reported that those agents couldn't carry out the obligation, but I find the data unconvincing in light of my qualitative look into OIC experiment 3.³⁴⁸ What appears to be people's judgments/intuitions in response to quantitative questions with a variety of available interpretations isn't satisfying to understanding if people find OIC intuitive for the very reason that what appears to be the case at face value isn't always so with scenarios of OIC. The second publication I'll introduce is by Moti Mizrahi and it is entitled "Ought, Can, and Presupposition: An Experimental Study." I regard it similarly to the first publication in that it contains a purely quantitative experiment that isn't sufficiently compelling for just the same reason when it finds that participants judged that an agent ought to do something that the agent couldn't do.³⁴⁹

The third work is in direct response to the Mizrahi piece I just discussed. In their reply, Kurthy and Lawford-Smith advance a number of valid criticisms of Mizrahi's study, two of which

are relevant to my current discussion in that they mirror similar worries I had toward the OIC study. First, they point out the problem of ambiguity in that “ought” is ambiguous in meaning, and it can be taken to mean multiple things by participants.³⁵⁰ I’ll elaborate on this complaint in the next section when I make just this observation in my own data (see discussion on definitional distinctions); while the scenarios Chituc et al. and Mizrahi use are different, they both leave “ought” to be interpreted ambiguously. This connects to a second, related point. Kurthy and Lawford-Smith assert “very few moral philosophers...think that ‘ought’ implies ‘can’ [*sic*] for every conceivable reading of ‘ought,’” which further iterates the importance of ensuring that participants are responding based on a desired reading of “ought.”³⁵¹ In short, we want to make sure empirical studies of OIC focus just on the meanings of “ought” we (most philosophers) expect to imply “can,” otherwise we aren’t finding anything of significance. And lastly, one thing they suggest in order to clarify how participants understand key concepts like “ought” is to include certain follow-up questions in an experimental design.³⁵² Though they don’t specify using a qualitative interview as I’ve done, it is clear my experimental design is aimed at being sure we understand how participants understood a given thought experiment and the concepts baked into it.

In response to the volley by Kurthy and Lawford-Smith, Mizrahi publishes the fourth work I want to discuss, “Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith.” Mizrahi takes seriously the given criticisms and runs an additional study that is taken to challenge, again, the intuitiveness of OIC.³⁵³ In my view, though, the measures Mizrahi takes are insufficient both in addressing Kurthy and Lawford-Smith’s worries and, more importantly, my MPEP worries. The purely quantitative surveys, vague enough to be interpreted in a host of ways, really don’t squash the concern over ambiguity and, based on my own study, I doubt participants

comprehended a newly added phrase “taking into consideration all the facts about this case” just in the way Mizrahi hopes³⁵⁴—if one added the think aloud method and follow-up interviews, my prediction would be something similar to what I encountered in my work, which is a variety of interpretations (more explored in the next section).

The fifth paper I want to discuss is a major work by Buckwalter and Turri entitled “Inability and Obligation in Moral Judgment.” Though it is expansive and thorough in many respects, none of its 8 purely quantitative experiments sufficiently account for the MPEP or the worries I introduce to Chituc et al.’s OIC study. In total, Buckwalter summarize their work as follows: “The main findings are that, contrary to prevailing philosophical wisdom, commonsense morality rejects the ‘ought implies can’ principle.”³⁵⁵ For reasons elaborated on in this total project, I find this conclusion to premature. Before moving to the next publication, I do want to elaborate on just one feature of the study which exposed to me an additional potential problem in much of the x-phi of OIC that I do not develop in full in the next section: just as I suggested in chapter two that compatibilism might be an incoherent position such that it might be unclear how to interpret participant responses in apparent support of it, the view that OIC isn’t the case might itself be incoherent such that it isn’t clear how we ought understand participant responses that seem to embody the view that OIC is false. I’m not here arguing that the view that OIC is false is incoherent, but rather pointing out that if it is, then it is difficult to know how we might understand participant responses that endorse such a view. In particular, though, I wonder if its potential incoherence, or perhaps the potential *obviousness* of OIC being true, might lead some participants to reason (explicitly or implicitly) that the questions they confront on surveys are *clearly* getting at something else—‘Surely, the researchers aren’t asking if someone ought to do what they can’t?’ This occurred to me when reading the way Buckwalter and Turri presented certain response

options on their surveys. For example, in one experiment, they asked participants about a man named Walter who promised to pick up his friend from the airport before being unable to do so due to car accident. The most selected description of his situation was that “Walter is obligated to pick up Brown at the airport, but Walter is not physically able to do so” with the other three choices being the other three possible iterations of Walter’s being or not being obligated and being or not being able to do so.³⁵⁶ On its face, this popular response might be taken to undermine OIC, but given the other available responses, it seems to best capture the vignette: there is an obligation in the story (regardless if it goes away), and Walter can’t make it.³⁵⁷ However, if OIC is obviously the case, or the falsity of OIC incoherent, it is entirely plausible participants pick this choice—the only choice that captures the obligation and the inability—without thinking it really possible they might be addressing something more nuanced and philosophical, such as the notion that they are endorsing that Walter is obligated to do the impossible.³⁵⁸ In reading the available responses, it might be obvious that no one ought to do what they can’t do, so participants might casually, or implicitly, take the question to clearly be asking something far more simple—i.e., whether Walter had an obligation in the scenario, and whether he can carry it out. In just looking at this one moment from the total study, it should be clear why the quantitative responses are subject to the MPEP; it seems impossible that the researchers could possibly know what participants meant in their responses to such vague and limited available responses.

The sixth paper I want to address is in direct response to this study by Buckwalter and Turri, and it warrants extensive summary here and further analysis in the next section. In essence, Kurthy, Lawford-Smith, and Sousa (hereafter Kurthy et al.), in their work “Does Ought Imply Can?”, treat Buckwalter and Turri’s work “Inability and Obligation in Moral Judgment” much like I treat Chituc et al.’s OIC: they worry about the experimental design, they re-create and modify it

in hopes to clarify how people are thinking about OIC, they include some qualitative data, and they include follow-up questions in a way that simulates a very brief kind of interview.³⁵⁹ And, in total, I was happy to find that Kurthy et al. and I largely corroborate each other's work in showing that participants' apparent negation of OIC is sometimes just the opposite once we are able to probe deeper into their reasoning, with the help of some qualitative evidence,³⁶⁰ such that we need to be careful in what conclusions we draw about laypeople's intuitions being unsupportive of OIC. Specifically, Kurthy et al. conclude that their findings from multiple studies that look at an agent's obligation to carry out an action in spite of some inability brought on outside of their conscious choices—i.e., unlike the high fault self-sabotage scenarios in the OIC study—“provide strong evidence that despite [Buckwalter and Turri's] claims to the contrary, people do make judgements largely compatible with the OIC principle, at least in cases in which the inability is not self-imposed.”³⁶¹ And what is of interest here and in the next section is 1) *how* they managed to reverse much of what Buckwalter and Turri found that is *similar* to what I captured in my study, and 2) *how* their study is relevantly different from mine.

To the first point, the crux of Kurthy et al.'s studies was clarifying if participants who ascribed obligation and inability to an agent were doing so in a way that really was in contradiction to OIC. For example, in their second experiment, they followed up with participants who selected that Walter (in a scenario modified from the Buckwalter and Turri vignette discussed above) is both obligated and unable to fulfil his promise of picking someone up from the airport with this question: “With this choice, do you mean that Walter is *still under the obligation* to pick up Brown at the airport *after* he becomes physically unable to do so?”³⁶² 74% of participants that faced this question chose “no,” showing that a response that initially seems to negate OIC actually supports it.³⁶³ This result should look familiar; in chapter six, I outlined a similar phenomenon with

participant 30 flipping her response on the “ought” question when asked about 5:55 (from a 4 that Brown “ought” to make the appointment to a disagreement with “ought” when asked about 5:55 in the interview), a time when (Chituc et al.’s) Brown could no longer make the meeting.³⁶⁴ That said, this reversal of Buckwalter and Turri’s finding also exposes a limit to Kurthy et al.’s study, that they acknowledge³⁶⁵, which is that they only test scenarios that look like Chituc et al.’s low fault scenarios wherein inability arises from events external to the conscious choices of an agent; they don’t ask about agents who self-sabotage on their ability to meet an obligation. And here’s where my study helps to complement their work, I did find evidence with participant 30, who worked with high fault scenario, of someone upholding OIC even when an agent self-sabotages.

To the second point, my studies diverges from that of Kurthy et al. in that I included much more (and more in-depth) qualitative data. Overall, I find their qualitative data to be corroborating of my findings in some ways, but also limited in others. To corroboration, they report some participant responses that really capture similar sentiments to what I found in looking at Chituc et al.’s study. For example, Kurthy et al. suggest that some participant responses highlight a “scope-shifting problem” such that some participants, rather than contradict OIC, viewed the scope of an agent’s obligation to change so as “to include new or alternate content.”³⁶⁶ With respect to a Walter scenario, one participant offered the following: “Walter made an agreement with full intention of keeping it and if he cannot fulfill the agreement, notice should be sent and a proxy should be appointed to carry out the agreement as specified.”³⁶⁷ This speaks to two of my conclusions (explored in the next section) that some people use “ought” to express there exists a moral remainder for agents to do the next best thing in situations where inability emerges, and that some people use “ought” to convey that one ought to do as much as possible to fulfil an obligation. However, I also find Kurthy et al.’s mix of mostly quantitative and some qualitative data limiting

and even subject to the MPEP for reasons I've outlined in this dissertation. While it is great to read participants essentially state that an agent has an obligation up to a point of inability—e.g., “He WAS obligated, but cannot physically do it so the obligation is no longer on him”³⁶⁸—at least some of such responses highlighted in the paper echo the language of the follow-up questions participants grapple with to such a high degree that I worry that participants' judgments, in being funneled through certain clarificatory questions, are being adulterated in some sense—e.g., the corresponding question that the quoted participant previously answered was, “With this choice, do you mean that Walter is *still under the obligation* to pick up Brown at the airport *after* he becomes physically unable to do so?”³⁶⁹ It is true that Kurthy et al. do address something like this and attempt to work around in a subsequent experiment, but I can't help but find the framing of the question in the fixed scenario somewhat suggestive and the lack of being able to clarify in long-form interview questions limiting. In the study run to remedy this biasing of language, participants are asked if an agent “is still obligated” which may plant the notion of *up to a point* in the mind of the participant.³⁷⁰ And even if all such language is corrected for, it is difficult to know precisely how participants are thinking about such scenarios without asking further probing questions for clarification. A strength of the Chituc et al. OIC experiment 3, and specifically my follow-up question about Brown at 5:55, is that there is no leading language in it.³⁷¹ When I ask about 5:55, I'm not (except maybe by exception) asking if Brown *still* ought to do something. And relatedly, my use of think aloud sessions is really helpful in avoiding biasing one response over another in that participants are reacting to non-suggestive questions as they encounter them before I seek clarification, which by its very nature might open the participant up to finding and reporting the response they think they ought to report. In total, though, this paper speaks not only to the value of combining more follow-up questions and more qualitative evidence to a study of people's

judgments, but it also unearths some findings that corroborate mine and are corroborated by mine. I will, then, include highlights from this publication in the next section as they become relevant.

The seventh publication, Leben's "In Defense of 'Ought Implies Can,'" is exceptionally similar to Kurthy et al.'s work. Leben takes up two worries about the studies from Chituc et al. and Buckwalter and Turri, namely the role of time in OIC scenarios and the polysemous nature of "ought," and runs two experiments using scenarios from Buckwalter and Turri such as the one about Walter promising to pick someone up from the airport. In the first experiment, he presents two groups of participants with an actual timeline of Walter's story (complete with images depicting the events), asking them either about his obligation in the clearly labeled time period *before* his incapacitating car accident on the way to the airport or the clearly labeled time *after* the accident and before the scheduled pick up time. Here he found participants' attributions of obligation before the accident were "significantly higher" than those from the participants who were asked about Walter after the accident.³⁷² This suggests that time-indexing is important for how people judge OIC scenarios. It also offers some kind of preservation of OIC since participants were given the opportunity to violate OIC in the condition after the accident and yet produced lower ascriptions of obligation than did participants when it is largely uncontroversial that Walter had an obligation: in the time before he became unable. In the second experiment, he, very similarly to Kurthy et al.'s experimentation, changes the language on scenarios from Buckwalter and Turri such that the available responses emphasize a relationship more revelatory of participant's thinking on OIC. For example, the structure that agent is obligated *but* unable then becomes the structure that agent is unable *but still* obligated.³⁷³ Leben, like Kurthy et al., finds increased participant support for responses that are consistent with OIC. Specifically, when compared to the original response structures, Leben's response structures produced significantly

different results in all three scenarios and a near reversal of support from obligated/unable (anti-OIC) to not-obligated/unable (OIC-compatible) in two of them.³⁷⁴

As with Kurthy et al.'s work, I found Leben's study to be helpful in corroborating some of my worries about how OIC is presented to participants—I also feel I corroborate their findings. In Leben's case, my emphasis on temporal concreteness is corroborated by both experiments in that they make an effort to ensure that “ought” is being temporally pit against “can”—without doing so I fear that OIC studies aren't *really* getting at whether people uphold or negate OIC. But again, like with Kurthy et al., I found Leben's work to be limited by the restrictions of quantitative survey questions. The timeline study presents participants *with an actual timeline* which might heavily bias them to think about obligation in terms of a structured timeline of events marked by periods of time before and after the onset of inability—i.e., the timeline says, ‘Think of the world this way, as before and after inability.’ The second experiment, with the altered language, forced participant responses through different descriptions of an agent's obligation and ability which means participant judgments are being funneled through choices that convey a kind of understanding of ought and ability—e.g., to read an option that says, essentially, that an agent is unable *but still* obligated is to take in a subtle kind of argument regarding the relationship between obligation and ability. This again speaks to the value of the think aloud session and interview, and my question about 5:55 in particular: while it references a specific time for participants to respond to, it doesn't “tell” them that obligation is understood along a timeline. It is up to *them* to connect this time to Brown's inability—we're back to my study here—and conclude if he is still obligated. Of course, any mention of time might bias participants to think about obligation in terms of time, but it is much less invasive to ask about a specific time than to funnel participant responses through a timeline complete with images of the events. In my interviews, participants could react to my

question about 5:55 open-endedly, perhaps even expressing confusion as to what 5:55 has to do with Brown's obligations.

The eighth work is Michael Hannon's "Intuitions, Reflective Judgments, and Experimental Philosophy," which, in looking at the merits and limitations of experimental philosophy³⁷⁵, focuses in part on Chituc et al.'s claims regarding OIC. In total, Hannon's work, which involves an experimental study based on the high fault condition of OIC experiment 1, nicely corroborates my finding explored in the previous chapter that participants sometimes maintain OIC by switching the temporal nature of the "ought" claim such that "ought to meet" is agreed with in an apparent effort to convey "ought to have met."³⁷⁶ It was good to corroborate Hannon's finding and have my work similarly corroborated, but, as is the case in all of the quantitative-heavy experiments, Hannon's work runs into the MPEP in that such a result isn't produced organically or offered up without probing from the participants (outside of an anecdotal class discussion). In my transcripts, I noticed that participants *did* switch the "ought've" language on their own accord, whereas Hannon's experiment offers that as a possible response in light of agent's situation. It is true that Hannon noticed this language switch in surveying students about the lunch-date cases (OIC experiment 1, high and low fault) and then discussing them. Specifically, students began switching to talk about the high fault scenario in terms like agent "*ought to have*," prompting similar utterances from other students, despite being asked if agent "ought to" keep the promise, all of which is a kind of helpful qualitative methodology.³⁷⁷ This I think is great in showing the value of a quantitative/qualitative feedback loop, though there are obvious worries (that Hannon seems to share) about making too much of student responses in a social, classroom setting. In the described discussion, what is important too is that more students began to recant their initial vote in violation of OIC (that agent ought to do what he can't) in the high fault scenario in favor of the

“ought to have” account; this activity was clearly helpful in revealing more complexity to people’s thinking regarding OIC. In his experimental study motivated by such thinking, though, participants were asked via Amazon Mechanical Turk about the high fault condition if they agreed that “it is still true”³⁷⁸ (just as OIC experiment 1 asked) that agent ought keep his promise at a time when he is unable to do so due to his blameworthy decision to wait around—70% agreed to Chituc et al.’s 60% greater than the neutral midpoint.³⁷⁹ This was followed by a forced choice question about the most apt description of the situation; this was the unique feature of this experiment. The most popular response among those who agreed with the first statement that agent “ought,” at 58%, was consistent with OIC: “At eleven forty-five, it is true that Adams ought to have met Brown at noon, even though it is now impossible for him to do this (because it would now take too long to get there on time). This is because Adams should have left earlier.”³⁸⁰ Only 34% chose the negation of OIC: “At eleven forty-five, it is still true that Adams ought to meet Brown at noon. It is impossible for him to do this (because it would now take too long to get there on time), but he is obligated to do what is now impossible.”³⁸¹ The third option, selected by just 8% of those who initially agreed, was the neither A nor B option; these participants were allowed to explain their responses, but Hannon makes no mention of whether they did.

At this point, you can probably spell out for yourself, reader, the problems I have with these survey questions! First, though, I do want to iterate the value in this experiment in that it gave participants the opportunity to negate OIC, albeit with leading language, and they didn’t popularly do so, showing that there really is more complexity with people’s understandings of OIC than might initially appear. However, yes, the force-question options look like they act as funnels through which judgments go, leading back to the MPEP and casting doubt on their accuracy in capturing unadulterated judgments. The negation of OIC response in particular, with “[Adams] is

obligated to do what is now impossible,” seems especially suggestive that there is something wrong in its selection.³⁸² Without conducting similar (and more controlled) focus-group discussions, interviews, or think aloud sessions, it becomes unclear if it is common for people to respond in terms of ought’ve and should’ve. Still, Hannon’s overall finding from his anecdotal evidence (which I surprisingly find most compelling despite the obvious potential issues with drawing firm conclusions from a class discussion!) and his experiment corroborate my finding that some participants switch the temporal reference of ought and should claims such that it seems they want to indicate not a violation of OIC, but something compatible with it. Recall that this came through even in the think aloud session for participant 30 (see chapter six)³⁸³, which highlights how natural and unprovoked the thinking can be.

The ninth and last work is by Yishai Cohen, entitled “An Analysis of Recent Empirical Data on ‘Ought’ Implies ‘Can.’” In total, it gives a helpful overview of the x-phi of OIC, and further iterates the findings of Hannon and Kurthy et al. that, more or less, preserve the intuitiveness of OIC in light of previous work suggesting a counter-intuitiveness to the principle.³⁸⁴ Again, in corroboration with my findings, this paper in large part emphasizes 1) the importance of temporality in determining whether an agent has an obligation or not and 2) how the subtle presentation of language—such as the word “still” in Chituc et al.’s experiment 1³⁸⁵—can influence temporally-based understandings of scenarios.³⁸⁶ There is good analysis here, but there’s no need to comment on it in any great detail since it clearly points to the same conclusion I’ve been iterating all along: basic quantitative surveys rely on an ideal story that there is reason to think easily breaks down. To close, though, I will highlight again that it is helpful to see how much emphasis Cohen, like previous researchers and myself, focus on the role of temporality with OIC (which I discuss more in the next section).

In sum, the literature in the x-phi of OIC has developed in line with similar worries as the ones I had in designing and executing my study: namely, that the findings showing people intuitively negate OIC are likely more complicated given the role of time and ambiguity in meaning of “ought.” The main works that initiated this literature, from Buckwalter and Turri and Chituc et al., have helpfully started a conversation about people’s nuanced judgments regarding unable agents and their obligations that later work has complicated in interesting ways. My findings—summed up as participants having messy judgments regarding OIC while often speaking in line with OIC—corroborate and are corroborated by some of these later works. The key thing missing from all of these studies though, is a move beyond quantitative surveys in any significant way, though some allowed for small amounts of qualitative expression. In this way, even though the later studies corroborate my findings, they themselves run into the MPEP from time to time as participant responses are capture via funnel-like questions that are potentially influencing participant responses. One of my unique contributions, then, has been to unearth people’s complicated OIC judgments more organically with fewer interferences. And this has also allowed me to identify more subtle contours of people’s thinking toward concepts related to OIC. In the next section, I offer my suggestions for future research and exploration based on my heavily qualitative research methods.

II. My Positive Contributions to the Debate

Rather than suggest a theory of OIC and attempt to pass some kind of burden of proof to other researchers, I’ll suggest seven things based on my findings, things that meaningfully connect to the x-phi of OIC. I do this for two reasons: 1) It is beyond the scope of this project to develop some narrow theory of obligation; 2) I collected people’s responses in a general fashion and I

didn't ask the kinds of targeted questions that could help "test" all of the many nuances of people's judgments regarding OIC. This process speaks to the value of triangulation of methods. Based on my suggestions here, researchers might conduct more quantitative surveys that account for some of the things I found, and then subsequently these quantitative data could be again explored in qualitative interviews. In my study, I let people talk quite open-endedly about issues relating to OIC as they connect to the vignettes by Chituc et al. I present these seven suggestions as seven paths down which future research could go.

ONE; IMAGINATIVE RESISTANCE OF A KIND MATTERS IN OIC SCENARIOS

My first suggestion is that scenarios designed to capture people's judgments relating to OIC need to account for a kind of imaginative resistance. In general terms, imaginative resistance "occurs when an otherwise competent imaginer finds it difficult to engage in some sort of prompted imaginative activity."³⁸⁷ I suggest this explains certain qualitative responses in the OIC study and my own study when participants appear to resist or reject the clearly described constraints of a vignette, such as the impossibility of an agent getting from one place to another in a given timeframe. I'm not concerned with the specific theory or kind of imaginative resistance such responses exemplify, and I want mostly to view the phenomenon in the broad terms I've just included.³⁸⁸ Without delving more into the literature on imaginative resistance, I conjecture that the theory I most clearly echo in my analysis of qualitative data below is in line with *Wontian theories* of imaginative resistance which "maintain that imaginative resistance can be traced to an *unwillingness* to engage in some sort of prompted imaginative activity: one simply *won't* imagine as one has been invited to."³⁸⁹ That is, it is not an inability to imagine that explains the resistance, something captured in *Cantian theories*³⁹⁰, but an unwillingness to do so, for whatever reason.

Future research, though, could connect my analysis below with more specific discussion in the literature on imaginative resistance. My main point here doesn't hinge on which kind of resistance we have in the data since I'm mostly interested in pointing out that such a resistance is apparent and influential on people's judgments related to OIC.

We see the kind of resistance of which I speak in experiment 2 of the OIC study and also in my own qualitative data; some participants don't seem to accept that an agent *really can't* make it somewhere in time. Recall in experiment 2 of the OIC study that Adams has made a promise to meet Brown at noon; the journey takes him 30 minutes but his car breaks down 15 minutes in.³⁹¹ Participants clearly read that Adams is in a lurch: "Because his car is not working at the time, Adams cannot meet Brown at noon, as promised."³⁹² And yet, some participants just weren't having it. When asked about Adams' predicament at 11:45, participants wrote the following to explain their quantitative responses: "At 11:45 he could still make it if he kicks his car in the right spot. He is not at fault for breaking a promise at this specific time"³⁹³; "Adams could take a taxi to his destination to meet Brown, and should do that or any other means of transportation because the meeting is important and he needs to remain professional"³⁹⁴; "Since he doesn't have to drive and park the car anymore, there is a possibility that he can still meet up with Brown at noon."³⁹⁵ These kinds of responses remind me of discussing the trolley problem with people unfamiliar with it when they deflect making a decision at first by asserting that they will simply shout to the workers on the tracks to move out of the way. This speaks to why a quantitative survey isn't always great for getting people's judgments; in an interview we could clarify to these participants and say that no, really, Adams *can't* make it. Then we see what they think based on acceptance of the constraints of the scenario.

In my study, I saw similar responses. Recall that I used vignettes from OIC experiment 3 wherein Brown made a promise to meet Adams and either 1) waited too long to make the appointment, or 2) had his vehicle break down such that he couldn't make it.³⁹⁶ From some think aloud sessions, I got the following responses: "Well, it says that it takes 15 minutes and it doesn't include any information that it perhaps if he drove faster he could make it in 10 minutes so I'll [unintelligible]..well...[sighs]...I mean it's he could make up a 5 minute difference...or he could...mm....hmm...[pause]...yea OK"³⁹⁷; and the following:

Uh..Brown is to blame for not making it to the theater by 6....uh...yes..well yes and no I guess yes because he was ready to go at 5:30 technically but then he decided it was too cold but even if it was too cold he could've just taken his car at 5:30..umm..and at 5:30 if he realized his car had broken down he could have also called his friend Adams who could have come and picked him up...umm.. or or he could have Ubered or he could have or if he realized his car was brok...not working he could have just rid...ridden bleh put on a thick jacket and taken his bike into the theater...so yea kinda...I mean the car breaking down isn't his fault...and can't really do anything about the weather but he kind of he really could have done a lot of other things....umm...so I'll give it a uh 4 'cause there're in circumstances he couldn't control.³⁹⁸

And the following:

K. At 5:50, Brown can't make it to the theater...uh do you agree or disagree with the following statements please circle your response. At 5:50, Brown can't make it to the movie theater by 6. Well it says that it takes 15 minutes to drive there and do all that stuff and it's 5:50 so I guess technically he wouldn't be able to make it but I think like with just 5 minutes like extra you can make up that time pretty easily if you just rush...umm...I

think he definitely could have made it but it's a little bit more iffy I think that he could have tried if he wanted to he could've made it but it doesn't seem like he necessarily wanted to so I'm gonna go ahead and disagree but only slightly.³⁹⁹

It's especially interesting to see how aware some of these participants are of the constraints detailed in the vignette and yet they still express some kind of resistance. They aren't, then, missing key information but rather trying to work around it.⁴⁰⁰

Of course, not everyone expressed such imaginative resistance to the confines of the scenario, but it is plausible that it nonetheless factored in people's thinking more often than was expressed. The reason we need to take this imaginative resistance seriously is that it might better explain why some people respond the way they do than a theory about OIC itself. If a participant doesn't *really* seem to accept that agent can't make a promised meeting, then it is difficult to say that she negates OIC if and when she judges that agent *ought* to make the meeting. In other words, we only find negation of OIC when someone clearly judges that an agent both can't and ought make a promised meeting; but if a participant's judgment of inability is flimsy, then it's unclear whether her "ought" judgment really is pit against the relevant inability in agent's situation. Of the two judgments—"can" and "ought"—it seems more important to OIC that the affirmation of inability is firm and unmoving than if the obligation is affirmed half-heartedly. In fact, even a half-hearted affirmation of obligation in the face of clear inability is suggestive of a negation of OIC, but not the other way around. In sum, I suggest that researchers need to be cautious of the subtle ways imaginative resistance might rear its head—whether for *Cantian* causes, *Wontian* causes, or other. And without the qualitative data from experiment 2 or my study, there's really no way we would've come to see this. And, I think my think aloud transcripts helped show just how clearly people were understanding the time constraints and still trying to get out of them.

TWO; FACTORS INFLUENCING OIC JUDGMENTS CAN WORK IN CONJUNCTION

Now that we have one factor on the table that might influence participant responses, namely imaginative resistance, I make my second suggestion: some factors that influence participants' judgments may overlap, intertwine, or generally work in conjunction with each other in ways such that they can't be fully pulled apart. This is something that other experimental philosophers have suggested as well; Nichols and Knobe, for example, suggest that a hybrid model accounting for multiple psychological processes might help explain the way people understand moral responsibility.⁴⁰¹ There are many examples of such intertwining that come up in the next suggestions I make, but I did want to highlight one potential overlap here in detail. Specifically, there is evidence that imaginative resistance and concreteness as I've discussed them so far work in conjunction to potentially explain participants' responses in interesting ways.

To show this, I want to make reference to experiments 1 and 2 from the OIC study as well as my own data, which emerged from my partial re-creation of experiment 3 in the OIC study. In experiment 1, an agent (Adams) promises to meet a friend at noon, a journey that will take him 30 minutes. In the high blame case, Adams waits until 11:45 at his home so he can't make the meeting; in the low blame case his car breaks down at 11:45 thus preventing the meeting.⁴⁰² Chituc et al. collect their participant response data on whether at 11:45 agent ought to make the noon meeting, but then they notice that some participants used the free-write space to suggest "alternative actions that Adams should have done instead, such as calling his friend."⁴⁰³ This indicates that maybe some participants weren't fully accepting that the noon appointment *couldn't* be made as such, and Chituc et al. notice that OIC supporters could claim that participants were conveying that Adams should meet Brown *even if* he shows up after the noon appointment time, not that Adams ought to

make the noon meeting time.⁴⁰⁴ Here we get our first sign of the role of imaginative resistance as I previously discussed it: some participants appear to not fully accept the constraints of the thought experiment and offer ways for agent to still meet an obligation that it is impossible to meet.

Chituc et al. notice the problem when participants offer alternative actions for Adams to take, so they conduct experiment 2 in part to make it clear that the noon cutoff is indeed a cutoff.⁴⁰⁵ In this vignette, if Adams misses the noon deadline, he isn't showing up to lunch late, but instead to the train station with no noon train to catch. And to make it crystal clear, they detail in the vignette that "since cell phones have not been invented yet, Adams has no way to contact [Brown]."⁴⁰⁶ Again, we see participants not fully accepting this (see examples above as I discuss imaginative resistance). In experiment 1, participants are told agent can't make the meeting, and yet they treat the 15-minute time window from 11:45 to 12:00 pm as something that *might* be overcome or worked around: "His car has broken down at that time. He could get a ride from his friend though"⁴⁰⁷; "Brown is probably already on his way, or at the restaurant waiting for Adams. Adams needs to get over it and go to lunch. As a courtesy, he should call Browns to let him know he's going to be late."⁴⁰⁸ In experiment 2, the train leaving and the non-existence of cell phones makes the 15-minute gap between 11:45 and 12:00 even more of a concrete cutoff, so to speak, and yet we see people still thinking agent can somehow make the meeting (see above examples). So this might lead us to conclude that some people judge certain actions as possible no matter what (or close to it): short of agent being struck and killed by a passing car, some participants might think there is *something* she might have done to make a promised meeting.

Here's where the last bit of evidence comes in, in this case from my study, to show that some imaginatively resistant participants might yet be convinced to attribute inability to an agent. It occurs when imaginative resistance is extinguished by a concrete scenario that seems sufficiently

restrictive. In other words, if agent is in a situation that lends itself poorly to imaginative resistance *and* it is made explicitly clear that agent's situation is being assessed at a very concrete moment in space and time, then participants abandon the alternative options for agent to take and do accept his inability. What makes this tricky to talk about with readers—like you, philosopher—who don't have this tendency to resist the confines of a thought experiment is that it seems odd that one action can be perceived as *more impossible* than another, or a deadline *more* of a deadline than another.⁴⁰⁹ The vignettes in the OIC study make it perfectly clear that the agents *can't* make their meetings; there is no uncertainty in this. Something similar occurs in the trolley problem: that agent can't shout to warn the workers of the trolley is to say that it is *just as* impossible to effectively shout such a warning as it would be for agent to freeze time and intervene—they are both equally impossible, which is to say impossible. And yet, it seems reasonable that when an action is made “off limits” or “impossible” by a vignette that would otherwise seem plausible—e.g., shouting at workers, grabbing a passing taxi—then participants understandably want to resist the “laws” of the thought experiment. Here's the point, then: what I found was what might be a kind of breaking point in people's imaginative resistance which occurs when the concreteness of agent's predicament was highlighted *and* the window opening to imaginative resistance was sufficiently closed. In sum, concreteness works hand in hand with imaginative resistance.

In this first example, participant 12 grapples with Brown's obligation in light of different moments in the scenario. What speaks to my point about the closing window on imaginative resistance is that when presented with 5:55, participant 12 at one point seems to leave some kind of possibility open to Brown getting there at 6, but at 5:59 she acknowledges the impossibility head on:

Me: And what did you choose for question 3 and why?

Participant: Umm..I said. . .umm. . .I was on the completely disagree side uh but it was a negative 2 umm. . .umm. . .and I think this one was a little harder for me ‘cause umm..it was saying that he should make it by 6 and he should ‘cause that’s what he had said but ummm. . .I guess if I was just answering this question. . .or not answering but choosing to how to. . .umm. . .I don’t know..agree or disagree with this I think I would actually be a little bit more on agree side just the wa..because of the way the question was stated or the statement was. . .umm. . .put but again like I don’t think anyone has. . .hmm..[whispers while saying “Brown ought to make it to the theater by 6”] . . .umm. . .well ahh can I change [laughs while speaking] that one again? Umm..I..uh I agree with with this one instead of disagree. . .umm. . .and I th.. ah so I had trouble with this because of what had happened I was thinking of like yea he should make it to the theater by 6 but also like things happened along the way and he couldn’t that he couldn’t control so I kind of disagreed with that statement but not thinking about that and just like from the beginning where him and Adams made the commitment yea he ought to make it to the theater by 6. . .umm. . .without knowing what else was happening.

Me: Great, to follow up on that, if we knew that Brown couldn’t make it. . .let’s say it’s 5:55 and he’s still on the side of the road still 10 minutes away from even getting there to park. . .so let’s say we knew he couldn’t make it would you agree or disagree with the statement. . .Brown ought to make it to the theater by 6.

Participant: Oh so if we knew all that. . .um. . .umm. . .no because I think..um. . .at that point it’s like it would it would it seems like he would be stressing to like make this appointment whereas. . .um. . .I mean he could re-schedule. . .umm. . .or call his friend and I’m sure that this situation is very understandable that his friend would understand. . .umm. . .but I mean if Brown was like at home at 5:55 and he was jus’ losing track of time at that point it’s like completely [laughs] his fault. . .umm. . .and he should find any means to get to the theater by 6 but knowing all that I think he. . .um. . .I think he has every right to tell that to his friend and explain not have to find a way to get there by in [laughs while speaking] like 5 minutes.

Me: And another follow up ‘cause you mentioned had he been at home just lost track of time. . .let’s say he decided he didn’t want to go and he just hangs around looks at this watch and sees it 5:59. At 5:59, in that case, would you agree or disagree that Brown ought to make it to the theater by 6.

Participant: Well that’s impossible [laughs slightly while speaking] at that point..umm. . .but ee..I think at that point he should still go to the theater. . .umm. . .and I don’t know and I guess be honest about it like what why he was really late. . .umm. . .but not blow off his friend just because he lost track of time.⁴¹⁰

In another example, participant 18 leaves open the possibility that Brown makes it at 5:50, which the vignette makes impossible, and then she expresses that at 5:55 or 5:59 this window is closed:

Me: And what did you choose for question 3 and why did you choose that response?

Participant: I put plus 5 so that I completely agree [sighs] I’m not sure if I understand understood this question properly meaning that I wasn’t sure if it meant that he should of made it like he should of just followed through with his promise with Adams and made it or if he could maybe drive faster or do something so that he could make it there I don’t know if that makes much sense but I put plus 5 because to me if you said you’re gonna meet your friend at the movie you should go. . .and it didn’t say that oh he didn’t tell he didn’t you know notify Adams like oh I’m not really feeling like watching the movie anymore he just sits down there and doesn’t say anything so what is Adams gonna do when his friend doesn’t show up for the movie so I think he should make it there so he should go and try and make it by 6.

Me: Great. Do you agree or disagree with the following statement? [participant says “OK”] At 5:55, Brown ought to make it to the theater by 6.

Participant: I disagree.

Me: So with that question you said you disagree and on question 3 which is similar but not exactly the same you put agree, what's the relationship in your responses there?

Participant: Well because 5 minutes have passed so then it's 5:50 and I think you have a better chance of making it in 10 minutes versus 5 minutes because perhaps you could make up for those 5 minutes when you should have left at 5:45 by driving a little faster or finding a parking spot something maybe even if you're standing in line and it's about to start you can ask the person in front of you can I go get this ticket something like that.

Me: And to bring it to one more follow up [participant says "Mm-hmm"], do you agree or disagree with the following statement? At 5:59 assuming that Brown hasn't left yet, Brown ought to make it to the theater by 6.

Participant: I don't think he'll make it to the theater by 6 if he..I would say though in theory he could go to the movie theater and ask to be let in but by then it would be past 6 and even though it is a strict policy exceptions are made all the time so if he really cared about Adams and about his promise he could still go at 5:59...he wouldn't make it by 6 but he'd have to put in the effort to go see the movie because of their policy so he'd really have to beg I guess.⁴¹¹

And in one more, participant 34 moves her responses from a 4 to a 0 in agreement that Brown ought to make it when the timestamp of 5:55 is introduced, and then she negates the question altogether when the window is closed to 5:59:

Me: What did question 3 on the survey ask?

Participant: Brown ought to make it to the theater by 6.

Me: What did you choose for question 3 and why did you choose that response?

Participant: I gave it a plus 4 because he ought to according to the plan that they had laid out...ummm...although I didn't put a plus 5 because there're other circumstances probably at play.

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: [pause] Mmmm...I probably say I I neither agree nor disagree because...uh..what he ought to do doesn't always match what he can do and I'm not sure if he it's even possible for him.

Me: Let's push it even further in time and assuming he has not left at this moment do you agree or disagree with the following statement, at 5:59 Brown ought to make it to the theater by 6.

Participant: I don't think it's an a question of whether he ought to or not I think it's whether he can or not.⁴¹²

These are important examples to digest because we might recall that experiments 1 and 2 in the OIC study had a concrete timestamp in the questions participants were asked, 11:45, which marked a time at which the agent could no longer possibly make the appointment.⁴¹³ It might have seemed as if Chituc et al. *did* account for temporal concreteness of a kind such that we could take their results to be unaffected by the concern that subjects aren't properly pitting an agent's inability with

his obligation when forming judgments regarding OIC. And yet, in these above examples, we see that not all concrete timestamps are treated equally by participants: what some participants see as possible in 10 minutes is different than what they see as possible in 5 minutes or 1 minute *even if* agent is said (in the vignette) to be equally unable in *all* such timeframes to make his promised meeting. What this suggests, then, is that concreteness alone can't account for some responses. For some participants, concreteness might only be a relevant factor for influencing people's judgments toward OIC if it is the kind of concreteness that closes the window for imaginative resistance.

This should hardly be surprising for the very reason that philosophers appear to design thought experiments to be *somewhat* physically realistic (even if they take place in odd realms) so as to make the actions available to an agent seem plausible. It is the imaginative construction of the trolley problem that in and of itself removes the possibility for some attempts to resist its confines. People know that a barreling trolley won't easily stop, they don't seem to doubt that being struck by one causes immediate death, and they can grasp the physics of switching tracks.⁴¹⁴ Check your own intuitions and conversations regarding the trolley problem here; do you find yourself or others hung up on these factors? I doubt it. All of these in effect close the window for imaginative resistance. We could construct a very different trolley-like problem with the same constraints on behavior while requiring subjects to take more of the philosopher's word on the topic, and I'd be surprised if people "bought it" the way they do the trolley problem. Here is my creation, the Marble Problem⁴¹⁵: "Imagine a marble is rolling down a road toward five people who don't see it coming. If the marble is not diverted from its current path, it will lodge itself under the shoe of one of the five people just as he steps down, causing him to slip and fall. If this happens, he will knock down the other four people as he falls and they will all five land awkwardly on their

necks, break their necks, and immediately die. You are tied up on the street lying face down just to the side of the marble's path and you can blow on it forcefully as it passes so it travels to the left of the five people where there is one person standing by himself. If you do so, the marble will lodge just under his shoe and he will slip and fall alone and break his neck and die. Because you are tied up and can only blow on the marble, and because only your strongest blow will affect the marble's trajectory—in just the way that it will divert it to the one person instead of the five—you can only do two things: blow on it as hard as you can so it causes the one person to die or not blow on it as hard as you can so it causes the five people to die. What should you do?" In other words, philosophers seem to appreciate the possibility of unrealism contributing to imaginative resistance when they opt for runaway trolleys over runaway marbles—not that there aren't plenty of weird or unrealistic thought experiments, such as those concerning zombies⁴¹⁶—and so we ought appreciate that for some people a window of 15-minutes seems to hold all the opportunity in the world while the same isn't true of a 5-minute or 1-minute window.

Accounting for different kinds of imaginative resistance might not be easy though. For one example from the surrounding literature, the researchers Kissinger-Knox, Aragon, and Mizrahi baked inability into their experiment by asking people about a person on a train who ostensibly ought to give up their seat to someone in need but is currently asleep.⁴¹⁷ In one sense, I think sleep is a perfect representation of inability since one is rarely in a kind of lucid dream state such that he or she might be able to rouse oneself to act. But at the same time, I can see people resisting viewing sleep as a state of categorical inability. While asleep on a train, apparently in an upright and seated position, someone might realistically be phasing in and out of a sleeplike state. And if there's enough wiggle room in such a scenario to allow for *some* kind of ability in the agent to act, then participants might ascribe ability to the agent in such a way so as to justify his or her obligation to

give up the seat. In total, it is unclear how resistant people are to imagining certain cases that might read as perfectly lucid to philosophers who design such scenarios.

THREE; CONCRETENESS MATTERS

By “concreteness” I mean those details of a scenario or a follow-up question that ask about specific, non-abstract details. For example, if someone is asked whether Brown ought to do x , then it is not as concrete as asking whether Brown, at time y , ought to do x . Concreteness isn’t solely about timestamps or temporal details, of course, because the very description one has about Brown can be more or less concrete based on how his appearance, values, disposition, goals, etc. are described. In short, my claim is that participants’ judgments related to OIC are affected by concreteness, and therefore we as philosophers should take this into account when attempting to understand how people think about OIC. If, for example, we learn that participants disagree with the statement “ought implies can” in the abstract (which we might view as the opposite of concrete) but continue to reify its truth in the face of concrete examples, then we’d be speaking inaccurately to report that they reject “ought implies can.” In the examples below, I analyze some cases where I see participants’ judgments being affected by or being sensitive to concreteness.

This focus of concreteness v. abstractness and its effects on laypeople’s moral judgments is not unique to me, a fact which strengthens my claim that concreteness matters. In their work on free will, Nichols and Knobe ask some participants about concrete questions of moral responsibility, such as whether a man in a deterministic universe who kills his family to be with his secretary is morally responsible, and ask other participants about abstract questions, such as whether someone can be morally responsible in a deterministic universe—interestingly, they find that participants tend to give compatibilist responses when presented with the concrete case

described and incompatibilist responses when presented with the abstract case described.⁴¹⁸ They later go on to grapple with how concreteness might explain aspects of how people reason about moral responsibility.⁴¹⁹ If we put aside the many important questions such experiments raise about the contents of a given concrete case, as well as the other factors that might be working in conjunction with concreteness, it is clear that concreteness is worth taking seriously when considering moral judgments or moral reasoning. From the same book, Sinnott-Armstrong has a chapter entitled “Abstract + Concrete = Paradox” wherein he notes that there are multiple areas and inquiries in philosophy, including some in ethics, that give way to some kind of paradoxical division between abstract and concrete thinking that might explain deep-seated debate.⁴²⁰ Perhaps features of the “ought” implies “can” debate are best explained along similar lines.

The salience of concreteness is further echoed in psychological research that examines the ways people are motivated to help victims that are depicted more or less concretely. Researchers Tehila Kogut and Ilana Ritov found that experimental participants contributed more money to help save an individual victim—i.e., a child described as victimized by disease—when they were given concrete details about them—i.e., name, age, and picture—and participants contributed less money to a similarly victimized child when no such identifying information was presented along with the description of their medical need.⁴²¹ Again, there are many more questions and issues worthy of exploration in their work, but for my purposes I want to note that this study too takes seriously the role of concreteness in moral judgment and reasoning. In both the work on free will and charitable giving, the increased concreteness of a case mattered in how people thought and responded. Similarly, I find in my study that concreteness matters for questions of OIC.

To start in on my data, here’s one example that builds on my previous analysis of the way imaginative resistance works in conjunction with concreteness. What follows is the entire

interview I had with participant 33 who originally selected 5 (Brown can't), 5 (Brown is to blame), and 5 (Brown ought); note that he offers language such as "macroscopic" and "microscopic" to highlight levels of concreteness in his thinking:

* * *

Interview Session

Me: OK I'd like to pick up with question 3 [participant says "OK"] to which you responded you agree that Brown ought to make it to the theater by 6 [participant says "Yea"] would you mind explaining your reasoning for question 3.

Participant: Sure so Brown ought to make it to the theater by 6 well first off uh ummm...when I read that I th...I ek see two potential meanings of of that Brown ought to make it to the theater by 6 the first meaning of it which I'm assuming wasn't meant was that given that he he leaves at 5:45 we can assume that he'll get there by 6 I figured that's not the meaning of the of the question...uhh...I figured the meaning of the question is...uhh... Brown has a responsibility to get there by 6 or Brown...umm...should get there by 6 or...umm...something to th..to th..that's what I assumed the meaning was so with that in mind Brown oughta make it to the theater by by 6...umm...I'm I'm thinking does I ask myself does he have like a responsibility to be there by 6 and then like you said my answer to that was ik...was I completely agree why do I agree with that statement that he has a responsibility to be there by 6 or he should be there by by 6 well first he he promised to meet his friend at the theater...he hasn't had a chance to see it but the last showing is at 6 o'clock that evening Brown is exc...and you know reading that I'm assu...I I guess I kind of assumed that they had agreed that they would be there at the same 6 o'clock eh uh ah showing I made that assumption so he promised that he would he would meet him there and...umm...I'm also assuming that this promise to meet his friends there was made with enough time to plan how he was gonna get there an an and such and how long it would take....so assuming that Brown had the time to plan when he was going to get there plan how he was going to get there and given that he promised to meet his friend there then...uh..I guess I kind of just conclude that he you know he has a responsibility to meet his friend there...uh...his friend has to get there also his friend's gonna be there so...umm...so yea yea based based on that based on the fact that he had enough time to plan out his trip there and to consider the various things that could go wrong...it's cold..you know...uhh...his ca...you know your car can have issues in the cold you could've planned for planned for that...umm..yea based on the fact that he promised and that he could've planned and should've planned things in advance I concluded that he ought to have been there by 6.

Me: Great so you used the phrase he has a responsibility [participant interrupts with "Yea"] and that's how you were thinking of ought [participant says "Yea"] in question 3. So let's say that at 5:55 his car and Brown are still he's on the side of the road and he's still 15 minutes from from actually getting there [participant says "Yea"]...so...umm...would you agree or disagree with the following statement, at 5:55 Brown ought to make it to the theater by 6.

Participant: At 5:55 he ought to make it by 6...umm...can you say the questio...can you uhh give the s...the scenario question one more time.

Me: Yea so if his car breaks down at 5:50 [participant says "Mm-hmm"] we're assuming now that it's been 5 minutes and he's still on the side of the road and [participant interrupts with "Mm-hmm"] there's there's no way he can get to the theater because his car's broken down and even more time has elapsed so it's 5:55 [participant says "Yea"] do you agree or disagree with the following statement, at 5:55 Brown ought to make it to the theater by 6.

Participant: So...I'm now kuh...uh...uh now there's a different dimension to this question that I'm that I'm that I'm...uhh...considering right now...so I think a mo...a moment ago I was looking at maybe I was looking at the situation...uh...from a kind of a macroscopic perspective kinda like in the grand scheme of things you know...umm...if I'm analyzing this story in the grand scheme of things he should ha...he should have been there on time but if I kinda zoom in an an and and and think like at this at this moment in time at 5:50 where his car is broken

down some unfortuitous...uhh...occurrence has happened...umm...at that moment I don't believe he has the responsibility to be there on time because now his car is broken down he's got this new situation to deal with...ummm...you know there's danger's involved with that so kinda like at this in this microscopic point of time or at this particular point of time for him in that moment...um...his concern needs to be his his safety and so when you m...when you said well at 5:55 does he have the responsibility to to be there by 6 eh...at that moment...um...his m...main...ummm...need...uhh...and I guess his greater responsibility's caring for his safety...uhh...taking care of his his you know his belongings his car an and making sure that that gets uh gets to a safe place I mean if it's cold outside is there snow I don't know is there are there weather elements that are making it unsafe for him to be there I guess what I'm getting at is at that at that moment in time his responsibility is is for his is for his safety so at tha moment of time you know ought he be there by 6...you know...no no but if I'm looking at the macroscopic perspective...um...I guess if I'm looking at like a macroscopic perspective I guess he he he should have been there by 6...lemme yea...I'm I'm seeing a different dimension to this questi... to this question...to this question now...umm...so...yea so if...ehm...if...umm...his OK so hi...ih he couldn't ha...maybe he couldn't have seen in in in advance that his car was was going to break down...um...I th I think I think that...uh...there's there actually there's a few different things going through through through through my mind now if we're talking about responsibility..ummm..so assuming tha... assuming that the the car wouldn't have have broken down...umm...he he di..he should've had the responsibility to be there by by 6 o'clock and only giving yourself literally 15 minutes that's not very responsible because lots of things can happen there can be someone in line more people in line there could be more traffic if it's cold there could be weather elements and then if you're not gonna get there right on time you're not gonna be able to go in so he should've he did have the responsibility to be there by 6 but once this unforeseeable thing you know it's not reasonable to assume that your car is gonna break break down nobody assumes that every day when they're making decisions so once the car broke down...um...I I kinda I guess that kind of negates the resp...th the responsibility he had to have been there by 6 o'clock so if I'm looking at that moment...uhh...I would disagree with the statement that he has to be there by 6...does that make sense? So at 5:55 and actually you know re-considering the situation at 5:50 I would disagree with the statement that he ought to make it to the theater by 6 but prior to that event happening yes he ought to make it there by 6.

Me: Excellent so in the sense that you read question 3 it sounds like you're describing when someone makes a promise they have a responsibility to follow through [participant says "Yea"] but once we get the specifics of the scenario and now there's a new situation that arises [participant interrupts with "Yea"] at this specific time he no longer can carry out [participant says "Yea"] make it so things change and so that responsibility [participant says "Yea"] changes.

Participant: Yea yea and I g..I guess previously when I was reading through it I didn't I didn't consider how I I wasn't considering enough of how the specific situation might...umm...might release him from that responsibility to get there get there by 6 ...yea... 'cause even you know even if he could have made it there by bike he would still need to take care of his car he can't just leave his car there you know so...ummm..yea.

Me: So when you were reading and answering question 3 [participant says "Yea"]...umm...were you considering a specific timestamp like at this moment so notice how I put the timestamp at 5:55 for [participant interrupts with "Yea"] that follow up question when you were answering question 3 were you thinking at a specific minute when when you provided your response?

Participant: I think...um...when I provided that when I provided that response I was...um...like I'm...umm...if if you had put if if it had said...umm...let me see let me re-read this question 'cause I wrote down 5:50ehh...OK I think originally when I was reading that question right there I don't think I was necessarily considering just the timestamp because if if instead it had said 5:55...umm..I eh if it had said 5:55 here I don't think I would've changed my responses 'cause he still probably couldn't have ma... he still couldn't have made it there by 6 and initially when I was the way I was thinking reading through these things I wasn't considering how he would have r...been r...released from responsibility and so I was I think I still would have been thinking that he should have been there by 6...does that [I interrupt with "Mm-hmm"] does that yea yea.

Me: I had a follow up and now it's it's slipping my mind so I'm gonna hold on and [laughs while speaking] try to remember what I just thought of [participant says "Sure sure"]. Let's imagine a slightly different scenario [participant says "Mm-hmm"] rather than Brown leaving at 5:45 and his car breaking down let's imagine that he decides I don't wanna go [participant says "Mm-hmm"] and so he stands around in his apartment for 10 minutes until it is 5:55 [participant says "Mm-hmm"] knowing that that will make it impossible for him [participant says "Mm-hmm"] to get

there [participant says “Mm-hmm”] so he’s intentionally closing the window on on being able to go [participant says “Yea”] but it’s not his car breaking down this was his decision to just wait and go well now I guess I can’t make it [participant says “Yea”] in that case would you agree or disagree with the following statement at 5:55 Brown ought to make it to the theater by 6.

Participant: Umm...I w...so..ehhwuh...uhh...the the immediate thing the immediate...uhh...statement that comes to mind is and and maybe this is just my my brain processing your question Brown should have made it to the theater by 6 I would completely agree with that statement...ummm...that he should have been there...ah..it’s obvious at that point that he can’t make it th..there by 6 because...uh...it’s it’s literally impossible for him to get there by 6 so should he get there by 6 at that moment well he can’t so...ummm...I can’t agree with the statement should he get there by 6 it’s more like should he have been there by 6..yes...ummm...Brown ought to make it to the theater by 6...uhh...he he can’t...I mean he still should he s...he still should have made it to the theater by 6 agr...and yah I still ag..agree with that statement but Brown ought to make it to the thea...uh...am I am I am I hashing the words too much here? [I interrupt with “No”] No.

Me [interrupting]: This is this [participant interrupts with “OK”] is good [participant says “Yea”]...so the should have is...ummm...as I’m understanding you he should have done the things that would have gotten him there by 6 because he made this promise.

Participant: Yea yea ou..or you know he should not have done those things which [laughs while speaking] which made it impossible for him to get there by 6.

Me: Perfect [participant says “Yea”]. To close, do you have any questions about the scenario or the 3 questions on the survey that you would like me to clarify?

Participant: Ummm...let me see here...anything about the scenario that I would like to have clarified...he can’t fix it himself in time to make it to....not that I can think of off hand no no.

Me: Perfect.⁴²²

* * *

This example is just to prime my more targeted examples to come. It seems clear here that concreteness was relevant to how participant 33 was thinking about concepts relevant to OIC: he expresses different judgments relating to obligation when considering a more general notion of Brown’s obligation than he does when considering it at time 5:55.

Importantly, this example from participant 33, along with one from participant 14 in the previous chapter, show that there is more complexity to something that Sinnott-Armstrong grapples with in his piece “‘Ought’ Conversationally Implies ‘Can.’” Specifically, they each point to an objection that he ultimately rejects to his case against the view that “ought” entails “can,” namely that a person is obligated up to the point at which meeting the obligation becomes impossible.⁴²³ My goal is not to fully refute Sinnott-Armstrong on his rejection of this objection,

but instead to challenge the main thrust of it which he grounds on intuitiveness—it is on the matter of intuitiveness that utterances from participant 33 and 14 become relevant. To start, Sinnott-Armstrong outlines a case in which an agent promises to meet someone at 6:00 pm but decides to attend a movie at 5:00 pm at a theater that is 65 minutes from the planned meeting location; this is the very case that the OIC study builds on.⁴²⁴ After 4:55 pm, then, the agent can't make the meeting. To this, Sinnott-Armstrong correctly points out that were “ought” to entail “can,” then it wouldn't be true at 5:00 pm or 6:00 pm that agent ought to meet for his appointment.⁴²⁵ But, according to Sinnott-Armstrong, “these implications are counterintuitive”—the promisee could reasonably tell the promisor that he “ought to be [there]” knowing full-well that the promisor can't.⁴²⁶ Again, Sinnott-Armstrong appeals to something like intuition (emphasis is mine): “[The promisee's] statement *seems* true, because [the promisor] did promise, the appointment was never mutually cancelled, and the obligation is not overridden.”⁴²⁷ Additionally, Sinnott-Armstrong argues that it “seems natural and true” that at 5:00 pm the promisee could tell the promisor that he “ought to meet [him] in an hour, and it takes more than an hour to get [there] from the theater.”⁴²⁸ That is, Sinnott-Armstrong has thrice referred to intuitiveness in defending this analysis. Then, he introduces that opponents have proposed that agent ought to meet an obligation up until he can't. His core rejection of this objection is a reiteration of the fact that he argues that it is true at 5:00 and 6:00 that the promisor ought to make the meeting, based on—as far as I can tell—his intuitions.⁴²⁹

And that's the rub: participant 33 and 14 both complicate the intuitiveness behind Sinnott-Armstrong's thinking (based on a thought experiment derived from the very one that Sinnott-Armstrong refers to). First, participant 33 initially agrees that Brown ought to meet at 6:00 pm despite his inability; this *seems* to align with Sinnott-Armstrong's intuition. It is true that

participant 33 is dealing with a low fault scenario whereas Sinnott-Armstrong is discussing something more like the high fault case in which the promisor chooses to make the meeting impossible. But that ultimately doesn't matter for this current discussion since Sinnott-Armstrong's reasoning isn't made on behalf of the blameworthiness in the promisor's actions (and later in the publication Sinnott-Armstrong introduces something like the low fault case and maintains that in it too the promisee could maintain that the promisor ought to be present for a meeting that the promisor couldn't make do to external forces rather than conscious choice⁴³⁰). The relevant moment from participant 33's explanation of question 3 is as follows:

I figured the meaning of the question is...uhh... Brown has a responsibility to get there by 6 or Brown...umm...should get there by 6 or...umm...something to th..to th..that's what I assumed the meaning was so with that in mind Brown oughta make it to the theater by by 6...umm...I'm I'm thinking does I ask myself does he have like a responsibility to be there by 6 and then like you said my answer to that was ik...was I completely agree why do I agree with that statement that he has a responsibility to be there by 6 or he should be there by by 6 well first he he promised to meet his friend at the theater.⁴³¹

However, when concreteness—temporal concreteness in particular—is introduced, it becomes clear that the participant 33 isn't upholding Sinnott-Armstrong's intuition. When asked if at 5:55 Brown ought to make it, he offers a long explanation which concludes with the following: "I would disagree with the statement that he has to be there by 6....does that make sense? So at 5:55 and actually you know re-considering the situation at 5:50 I would disagree with the statement that he ought to make it to the theater by 6 but prior to that event happening yes he ought to make it there by 6."⁴³² Here he reports something like the objection Sinnott-Armstrong rejects. Before the car breaks down at 5:50 (which is presumably what participant 33 means by "that event"⁴³³),

participant 33 agrees that Brown ought to make it; after it becomes impossible he disagrees that he ought to. This looks very similar to the notion that the obligation holds until the moment that ability is negated. What this suggests, then, is that Sinnott-Armstrong's intuition might not be *incorrect*, but instead it might be either *non-universal* or perhaps *incomplete*. I think the latter is most plausible. Some general sense of "ought" might hold for some impossible actions, but when we zoom into a concrete case, some people flip their judgments—a complete view tells this whole story and perhaps even more complexity. Therefore, the question of entailment as pertains to OIC might not be so straightforward and easily grounded on the intuitions that Sinnott-Armstrong articulates.⁴³⁴

With participant 14, who worked with the high fault scenario, we get another example of someone ascribing obligation up until the moment of inability. During her think aloud session, participant 14 reacts to question 3 in the following way:

Brown ought to make it to the theater by 6. So I feel like you can't really contemplate that unless because of the the way the verb tense is you you can't contemplate unless it's like if up to 5:45 if it's after that an..then it's already too late but...umm...I...hmm...I'm th...I'm thinking like it's a crappy thing to do to your friend because your friend has now spent the money and is seeing this movie and a lot of people don't like to see movies alone.⁴³⁵

It seems clear that her response is somehow contingent upon if she is being asked about the time before 5:45, when the obligation is still possible, or after 5:45 when it no longer is. She iterates that one can only address the question as phrased up until the breaking point, but then after it isn't clear if one can respond to it as it is phrased; after 5:45 she seems to indicate that the question loses meaning or clarity because Brown can no longer meet his obligation. Like participant 33,

participant 14 doesn't share in Sinnott-Armstrong's intuition that an "ought" remains even after an obligation is impossible to meet. Were she to share in his thinking, she wouldn't have placed 5:45 as the breaking point for ascribing "ought"—she'd have been happy to contend that even at 5:55 that Brown ought to make it. In the interview, when asked about 5:55, participant 14 even offers this exchange which shows her reluctance to ascribe obligation once ability has been negated:

Me: Great. Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: [laughs] So I'm...I don't think I c..do I have to agree or disagree..I don't think I can answer that because that also like on a grammatical level for me implies that he's on the way and the question is like is he gonna make it and like yea you know he he should be able to make it that that's how I hear that question...[I say "OK"]..can you read it again?

Me: At 5:55, Brown ought to make it to the theater by 6.

Participant: I don't think that we have enough information to answer that...or I have enough infor..[unintelligible]...the way I'm conceiving of it the question and the situation I don't think I have enough information to answer that.

Me: If we knew that at 5:55 Brown is still at home...

Participant: Then no..he can't. [I say "He can't..."] He can't make it by 6 because if it's true that it takes 15 minutes then then no.⁴³⁶

In this dialogue, participant 14 shows no eagerness to endorse the intuition Sinnott-Armstrong has, that the promisor ought to make the meeting despite his inability to do so. Instead, she seems so tied to the need for obligation to be coupled with ability that when asked if Brown *ought* to make it, she eventually responds that she disagrees that he *can*—the "ought" is used seemingly interchangeably with "can." In total, this points to the significance of concreteness in understanding concepts of obligation and ability. When concreteness, especially temporal concreteness, is introduced, we gain more insight into precisely how people understand these OIC scenarios. Of course, it seems that Sinnott-Armstrong appreciates the concreteness of the situation he analyzes and argues nonetheless that the promisor ought to make the meeting *even* at the moment of the agreed up on time, 6:00 pm, when the promisor and promisee are some distance

away. If this is intuitive to Sinnott-Armstrong upon full appreciation of concreteness, then it isn't to participants 33 and 14 and potentially many other people, all of which suggests his argument's reliance on intuitiveness isn't so sturdy. And, more importantly, it suggests that when Chituc et al. see participant responses that *appear* to coincide with a certain intuition, they need to push on it to see if it is really a completely different judgment underlying it. Concreteness might not matter to the trained philosopher who *naturally* appreciates its role in a thought experiment, but it seems to matter significantly to some laypeople when they engage with the scenarios.

Before looking to more transcripts, I want to briefly reiterate Leben's timeline experiment to make it clearer how it corroborates my finding that temporal concreteness is relevant to how some participants understand an agent's obligation. Leben, who presented participants with a timeline of a similar scenario to the low fault condition of OIC experiment 3, one in which agent is incapacitated after a car accident, found lower ascriptions of obligation after the accident than before.⁴³⁷ This emphasizes how important it is to capture participants' thoughts on OIC as contextualized by temporal concreteness since agreeing that an agent has an obligation takes on very different meanings if that obligation is tied to a time period of able action or one of unable action. At the same time, my study helps corroborate and push beyond Leben's work in that I unearthed talk of dividing an agent's obligation into periods separated by inability *without* showing participants a detailed timeline; this shows that my participants' temporally-based responses weren't so primed by a hint that they *ought* think in such temporal terms—a hint made manifest in the depicted timeline. My interview question about Brown at 5:55 certainly introduces a temporal element into the conversation, but it doesn't force a response *on behalf* of a pre-ordained timeline. This is the value of qualitative methods: they allow for more organic responses.

The next example I'll analyze is from the interview with participant 31, who completed the low fault version and responded 5 (Brown can't), -4 (Brown is to blame), and 5 (Brown ought). Here is a relevant exchange from her interview that highlights temporal concreteness specifically:

Me: What did you choose for question 3 and why did you choose that response?

Participant: Umm...I choose uh I completely agree because he should have...umm...he should've been able to make it by 6 but he was waiting too long.

Me: When you say he should have been able wh...uh..do you mind elaborating on that what do you mean by he should have been able to make it?

Participant: So like if he would've left at like...he so it says that he...umm...when he drives his car it would take 15 minutes and then if he would ride his bike it would take 30 minutes so if he tried to ride his bike eh...er if he tried to start his car at 5:30 even if it stopped working 5 minutes later he would've had enough time to take the bike.

Me: So in other words he could have left at a different time and that would've [participant interrupts with "Yea"] ensured that he made it [participant says "Yea"]...do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: I th..yea that's I completely...he would not make it.

Me: So would you agree or disagree with that statement?

Participant: He wou..it's the statement is that he would make it by 6?

Me: At 5:55, Brown ought to make it to the theater by 6.

Participant: So I disagree because he wouldn't be able to make it in time.

Me: So with the question 3 you said you agree that he ought to make it by 6 that [participant interrupts with "Right"] doesn't specify what time [participant interrupts with "Mm-hmm"] we're talking about, when I'm talking about 5:55 you said you would disagree [participant interrupts with "Right"] do you mind explaining the [participant talks over my next word with what sounds like "The difference" but it is somewhat unintelligible] connection? Yea.

Participant: OK so...ummm...when it says Brown ought to make it to the theater by 6 in my head I'm just thinking it of like a general assumption someone ought to be able to make it at that time like if I had work and yea I ended my job at 5 I would tell my friend that I ought to be at home by 6...but like if it's like a time thing then you wouldn't make it in that timeframe.⁴³⁸

The first noteworthy moment here is that despite selecting on the survey that Brown can't and ought to make it, participant 31 disagrees that at 5:55 Brown ought to make for the very reason that he can't do so. That is, when there was a concrete time on Brown's "ought," namely 5:55, she changes her answer to uphold "ought implies can" whereas without the concrete time, she

appeared to negate OIC. The second noteworthy moment is when I ask participant 31 about this very discrepancy, she articulates something just as I am arguing: that without the concrete timestamp she views “ought” as some kind of general (or I argue “abstract”) sentiment as to what one “ought” to do, but the concrete time constrains this “ought.” I take this as evidence that her survey selections merely gave the appearance of negating OIC because nothing in the two questions pit “ought” and “can” against each other—i.e., concretely—in a way that would be relevant to evaluating OIC. This tells me that concreteness, and specifically temporal concreteness is relevant to understanding OIC.

Next is an excerpt from the interview with participant 15 who responded to the survey with 5 (Brown can’t), 0 (Brown is to blame), and 2 (Brown ought).

Me: And what did you choose for question 3 and why?

Participant: Umm..I chose plus 2 agree because...ummm... I think it makes sense to leave on time but I don’t completely agree that he should make it there by 6 because things do happen and do make you like 30 seconds or a minute late sometimes.

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: I I disagree. [I say “Would you mind elaborating on why?”] Umm..yes well my first thought was right now where he is you know in sort of his route to get there could he make it...ummm...and still no because he’s still too far away...ummm...and so it doesn’t really matter why type of transportation he takes he’s not gonna make it there.

Me: So...initially when you answered question 3 which does not specify at 5:55 [participant says “Mm-hmm”] you said you agreed that he ought to make it [participant says “Right”] and then when I incorporated at 5:55 you said he would disagree...what’s the relationship between those responses?

Participant: Umm...well I guess his...what he ought to do I guess like his his path is..eh uh..already sort of decided for him at 5:55 he already knows that he’s gonna be late and that he doesn’t have any way to get there...ummm...but when I read question 3 I assumed that it was talking about before he’d even left so ideally he would have gotten there by 6.

Me: And before he would have left like when he was considering riding his bike at 5:30 or...

Participant [interrupting and talking over my last word in the previous statement]: Exactly.⁴³⁹

Again, this participant produces survey responses that indicate, on their face, a negation of OIC given that she affirms Brown can’t and ought to make it. And again, when presented with a concrete time for whether or not Brown ought to make it, 5:55, this participant disagrees that

Brown ought to make it for the very reason that he can't. But, interestingly, *how* this concreteness affects her answer appears to be different than with participant 31. For participant 31, the concreteness appeared to take her general judgment of "ought," which wasn't tied to a specific timeframe in Brown's journey, and bring it to a specific time in Brown's journey, namely 5:55, thus leaving her to disagree that he ought to make it. But for participant 15, the concreteness appears to cause the timeframe being considered to actually change from one timeframe to another. Participant 15, rather than affirming that she previously thought of "ought" as applying generally, expresses that she was imagining the "ought" pertained to some time before Brown left for the appointment. In both cases, temporal concreteness sharpens up what the question of "ought" applies to, leaving the participants to affirm OIC rather than deny it.

My interest in the role of concreteness in people's judgments relating to OIC extends beyond temporal specificity. In the excerpts below, I show that another, more general concreteness is apparent in that people steer *away* from concrete examples when asked if there are things that people ought to do that they can't do. Just as providing concrete cases in which "ought" and "can" are pit against each other leads some people to affirm OIC, some participants had difficulty negating OIC with anything like a concrete example.

Participant 2 completed the low fault version and provided the following response when asked directly about OIC:

Me: Are there things that a person ought to do that he can't do? Please explain.

Participant: Ummm...that's a really good question...ummm....Guess the first that comes to mind is people or a person ought to make moral decisions..ummm...based on...ummm...my own personal beliefs but if they weren't ... um...raised in the same type of house or if they didn't have the same values and upbringing...they possibly can't make those same decisions... 'cause they..they're not capable or the capacity of of understanding is sometimes different you know, like ummm...you know if you're like a new immigrant to the country as opposed to you know you're a fourth or fifth generation college graduate you see things completely different, for a fourth or fifth generation college graduate could say something to an immigrant that would be like they ought to know better in sincerity they really don't...you know.. so it..it can be confusing at times I guess.⁴⁴⁰

Here, when given the opportunity to conjure any example in which ought doesn't imply can, participant 2 offers a great example of the kind of things people ought to do (i.e., act morally) while connecting it to a kind of epistemic limitation, that sometimes people don't know what actions would be moral in a given context. On the surface, we have a moral "ought" that is prescribed to people who "can't" carry out this "ought"—this might be taken as a negation of OIC. And yet it seems clear that this general notion of obligation is not in any real conflict with the mentioned inability for any specific or concrete scenario. The example seems to better exemplify the sentiment that there are things people generally ought to do but sometimes people's upbringings make it difficult for them to do so. It hardly reads like the claim that at a time t , agent x is definitively incapable of doing y as a result of their upbringing and at time t agent is obligated to do y despite it being an impossibility.

Next, participant 6 said the following when he was asked about OIC (please ignore the first words from a previous bit of dialogue):

Me: Unforeseen circumstances [participant says "Yes, yes, well said"]. Are there things a person ought to do that he or she can't do? Please explain. [participant says "One more time please"] Are there things a person ought to do that he or she can't do?

Participant: Umm..[pause]..I think in..eh..the broad sense of like that question is like there are a lot of things that people ought to do in life I mean we ought to do we ought to be better people we ought to you know provide for family whatever it is we ought to do a lot of things..we ought to get off our asses and go work..uhh..we ought to go hit the gym..we ought to eat well..umm..but there's a lot of reasons that people don't do what they do and that's either a personal issue which in which they actually have no drive to do those things they just don't want to do them..um..and you said that that they can't do [I say "Mm-hmm"]..[unintelligible]..yea and there's also things like that so I mean there's also unforeseen circumstances, then there's also people have uhh mental illnesses or physical ailments that kind of can prevent them from doing those things...umm..yea...yea.⁴⁴¹

Similar to participant 2, participant 6 addresses the very general kind of "ought" that people should strive for, from self-care behavior to moral behavior. Though he notes there are reasons people don't and can't do these general things. But in no case does he give a concrete example of a person who is physically incapable of doing something at time t being viewed as obligated to do that impossible thing. The sentiment is much more the following: there are things that are generally

good for people to do; some people don't do them, and some people can't. These hardly strike me as statements undermining OIC precisely because they lack concreteness.

This too is something that Sinnott-Armstrong grapples with in his defense of "ought" as conversationally implying "can." Specifically, he addresses an argument for moral relativism which articulates, in essence, that because someone brought up in a specific culture can't imagine doing some action x that is deemed good by one's own culture, such as caring for one's children, then they oughtn't do x .⁴⁴² Importantly, Sinnott-Armstrong notes that OIC is a premise in this argument such that if OIC is the case, then it seems as though moral relativism wins out. Sinnott-Armstrong suggests that rather than allowing for such relativism, we might, on the basis of "ought" conversationally implying "can," find it meaningful to say that people from this other culture "ought to care for their children, but there is no point in saying so to their faces, and they are not blameworthy for not doing so."⁴⁴³ Indeed, this is a viable solution if it is the case that "ought" only conversationally implies "can" but doesn't entail it. But from certain participant responses, which I discuss below, there is reason to think that this might be counterintuitive or at least not a full representation of people's thinking. Of course, I didn't ask people about cultural differences and moral relativism, but the above response indicates that people might allow for less relativism when talking *generally* about moral actions and less so when discussing *concrete* ones. This again emphasizes the role of concreteness in how people think about moral behavior and obligation. To elaborate, let's look at the above example as well as an excerpt from participant 14's interview.

Just above, participant 6 is directly asked if there are things a person ought to do but can't; he offers up something quite general that seems to negate moral relativism in its unqualified generality, at least on its face. Specifically, people ought to do certain general things, he says, like "provide for family" and "eat well."⁴⁴⁴ But earlier in the interview, in response to a very concrete

scenario wherein someone can't do what they ought to do, participant 6 (who originally selects 5 that Brown can't and 5 that Brown ought make the meeting), ends up agreeing with something like OIC:

Me: Do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: Ummm...[pause]...I uh..I would disagree with that because umm given the scenario and given that his car broke down I would say that it's still ideal that he makes it there but at 5:55 I think he should definitely tell Adams I'm not gonna make it.

Me: Um..so..when you answered question 3 you had said that you agree and then for that last question where I asked about at 5:55 you said you disagree [participants says "Mm-hmm"], what's the relationship between your responses there..what do you see?

Participant: Yea, so of the kind of the difference that I see in that is umm..I see that at at 5:55 I'm already given the..you know..when I'm reading question 3 Brown ought to make it by 6 I'm assuming..um..with that I'm assuming that the question kind of gets me to it's the point of..uhh..it's still 5:30 or it's still earlier when we made the plans umm..so I'm thinking at that point you should definitely still make it whereas at fif..5:55 I'm now understanding that you know something happened and at 5:55 to make it by 6 seems you know almost impossible basically.⁴⁴⁵

The upshot is this: what one ought to do in general might apply more universally (a kind of negation of relativism) in a way that seems to negate OIC, but then OIC might obtain when it is applied to specific cases (which could support the relativism argument). Note that participant 6 even says that in "it's still ideal"⁴⁴⁶ that Brown makes it at 5:55, but that he disagrees that he ought to.

And participant 14 offers this response when asked the direct question about OIC:

Me: OK, perfect. Are there things a person ought to do that he or she can't do? Please explain.

Participant: Hmm..[pause]..well..it raises the question of whether you can whether it's legitimate to say that someone ought to do something that they're incapable of...um...I I think people might perceive that someone ought to do something that they're incapable of doing...umm...if it's physically impossible then I don't think it's reasonable to say that they ought to do it...if i...I'm thinking of like if you have the emotional resources or the energy to do something and n you really might be struggling and so maybe there is something that you should do but you make a choice that you're...you don't feel in the moment able to do it and that that feels like those things can kind of be in tension but if it's literally impossible like if someone is paralyzed and they see a person who's drowning or whatever it you know like we could say a human should try to save a drowning person but that person's not capable and so I yea I'm not sure whether that person ought to then...it's just like a generic person ought to do this but if you if you don't have the ability to do it it's not your fault [laughs].⁴⁴⁷

Again we see general claims that seem to negate OIC and a kind of moral relativism—i.e., “a human”⁴⁴⁸ ought to attempt to save someone from drowning—but then a reluctance to negate OIC

and moral relativism when someone is paralyzed and physically can't save the person drowning. Of course, we might get very different responses if we introduced the notion of what people from another culture ought to do. But the point is that what might be going on in what Sinnott-Armstrong is grappling with isn't that people are sharing in his endorsement of "ought" conversationally implying "can" and bypassing relativism; instead, people might really endorse a relativism when individual concrete cases are explored but endorse more universal norms when dealing with general notions of what people ought and ought not do. This all speaks to the complexity of people's judgments regarding obligation and morality and the need to account for abstractness and concreteness in understanding such judgements.

Let's move on from this point about relativism and explore one last example wherein concreteness seems relevant. Participant 11 had the following to say when asked in the interview directly about OIC:

Participant: Umm..just like anything [laughs]? [I say "Yea"] Uhh..sa..ask the question again sorry [I say "Are there things a person ought to do that he or she can't do?"]...mmm...ought to do that they can't do...yea I mean it's such a vague question..yea I guess.

Me: Is there anything that comes to mind, an example?

Participant: Umm...I don't know..I guess the first thing that came to mind was like family obligations...um...like you you should do those things but sometimes you have other plans...or you know...something else sounds more exciting and so [laughs while speaking] you say no and you can't do it but you should do it I don't know...so that was the first thing that came to my mind.⁴⁴⁹

Again, the prescriptive "ought" is some general sentiment of what someone should do, not a claim about a concrete inability being pit against a concrete obligation in just the moment of that inability.

Whether asked directly about a concrete time in the Brown scenario (5:55) or asked to conjure a counterexample to OIC, some participants respond in ways that highlight the significance of concreteness. This tells me that future research needs to account for this factor before concluding that an agreement with an "ought" statement can be taken at face value.

FOUR; DEFINITIONAL DISTINCTIONS MATTER

In my qualitative data, people appear to judge Brown's actions based on a variety of different definitions of "ought," which is what I am referring to as the finding of "definitional distinctions." If people, laypeople and philosophers alike, happen to view "ought" in different ways in relevantly similar contexts, then I think this should tell philosophers to tread more carefully in how we analyze and discuss OIC.

The first few excerpts point to one interesting way some subjects appeared to understand "ought," which is something like an obligation to make a full effort despite obstacles. Consider participant 28's response to a question about OIC.

Me: Can you think of an example of when you might say someone ought to do something even if they couldn't?

Participant: Mmm....like if someone was...so say if someone were to like regularly go visit their grandparent somewhere to help them around the house and like they're like physically injured or something I still think that they should go just in order to stay with the commitment that they make.⁴⁵⁰

In this moment, we learn that perhaps her use of "ought" isn't the kind that would negate the principle OIC: here she treats "ought" as doing *what one can do* despite physical limitations. This notion of obligation, doing *the most one can* is actually in line with OIC, not a negation of it. A negation would affirm that the physically injured person ought to do what her physical injury prevents her from doing. Kurthy et al.'s work nicely corroborates this understanding of obligation; they report multiple instances of what they deem the "scope-shifting problem" wherein subjects shift the scope of an obligation from the original obligation to some sense of trying to do something in light of an inability to carry out the obligation as planned.⁴⁵¹ For example, in response to a vignette about a playground safety worker tasked with protecting children on a playground from broken glass before becoming stricken by paralysis, one participant responded that the agent is obligated to remove the glass on the following explanation: "In some way if he knows there's

broken glass and no one else is notified, there needs to be a way he can communicate with someone he can or warn the kids about it.”⁴⁵² That is, he can’t do the original task, but there’s still *something* it seems he’s obligated to do, just as participant 28 suggests the injured person still can do *something* in light on an obligation, namely show up.

Similar is participant 38, who completed the high fault version with 4 (Brown can’t), 5 (Brown is to blame), 4 (Brown ought). He had this to say about his response to question 3 (on ought):

Me: What did you choose for question 3 and why did you choose that response?

Participant: Uhh...plus 4 nearly completely agree...umm...‘cause Brown he did promise Adams that he’ll meet him there so he sh...should have to do his best to get there by the time that they decided...umm...even if Brown doesn’t wanna go but then Brown should call Adams just tell him he doesn’t wanna go instead of I guess just making Adams wait there for him...umm...and maybe Adams also misses the movie then ‘cause he’s waiting outside for his friend Brown...uhh...the only thing that held me back from saying completely agree is that Brown is free to make his own decisions and if he wants to be kind of mean to his friend he has every right to do so like he’s no obligation to to make it on time...umm..but it’s like the right thing to do like he he promised he’s his...like they’re both friends...umm...I dunno I guess that’s friends do like yea.⁴⁵³

Here we have survey data that appears to negate OIC in that participant 38 affirms Brown can’t and ought make it. And yet, participant 38 reveals that his reasoning for selecting ought had to do with Brown making his best effort, which isn’t the definition of “ought” that would be relevant for undermining OIC since it is based on the acceptance that one *does their best given certain restrictions*—i.e., one does the best that they *can* do, not what they can’t do.

Another way “ought” was understood was in relation to the severity of Brown’s situation. During the think aloud session, participant 8 had this to say in reaction to question 3 (about ought):

OK, number 3. Brown ought to make it to the theater by 6. He ought to make it...he should make it to the theater by 6...umm...that seems logical because his movie starts at 6 and you want to be there in time for the movie since the theater won’t let you in if you arrive anytime after 6. But in terms of getting sort of zooming out and looking at it you should make it to the theater by 6...of course you should but really..[unintelligible]..this is where

we get into the philosophical part of things because in order to see the movie you should make it there by 6 but it's not like you know this huge requirement if something comes up and you can't it's gonna be OK...so..I'm gonna say that I agree with that because in his current scenario he wants to see the new movie and he's excited about it so in his mind this is probably the right choice but I'm gonna go with like a plus umm...3 because we have to account for potential other things happening too. OK.⁴⁵⁴

She indicates that the stakes just aren't so high, and this seems to factor into her response, or potentially does. While participant 8 doesn't explicitly define "ought" as pertaining only to serious matters,⁴⁵⁵ it appears that her treatment of "ought" is to some degree grounded in the severity of the consequences of the scenario, which are low.

Similarly, participant 12 said the following in her think aloud session in response to question 3:

Brown ought to make it to the theater by 6. I mean it's just a movie. Um...[pause]..and...umm...I think yea it's an appointment that he it's a commitment that he'd made with his friend...um..and he should of strived to get there by 6 but considering the situation it ..it seems like it was out of his control...um...he did leave right on time he s'didn't leave earlier and that's OK 'cause he'd calculated for to get there by 6 and that would've worked out..umm..hmm...[whispers unintelligibly]...umm...hmm...umm..I think th...Brown ought to make I think the wording of this like statement is making me feel like....umm...like it's like he...I don't know it's..it's a social obligation not like a professional thing it's not something that...um....he's relying on for I know like his basic like his necessities...umm...ummm...so I kinda disagree but not completely so ne..do negative 2 for that, yea.⁴⁵⁶

This example reiterates that there is an important definitional distinction that needs to be accounted for wherein “ought” might contain some reference to the severity of consequences of a situation.

In the case of participant 36, we see an explicit mention of multiple meanings of “ought”; she says the following during the think aloud session when encountering question 3:

Brown ought to make it to the theater by 6. I’m not 100% sure...ummm..on the meaning of the statement if the statement means Brown should make it in time logistically I would say no I disagree and if it means Brown should [laughs while speaking] make it there from like a [laughs while speaking] an ethical perspective because he promised to meet his friend there...ummm...I would say I completely agree with that so I’m not sure which..ch..Brown ought to make it to the theater by 6 out of respect for his friend completely agree. Is that what the statement was intended to..logistically no he won’t make it there by 6 if he left at 5:50.⁴⁵⁷

What’s obvious is that it’s *not* obvious to participant 36 how to interpret question 3. What isn’t obvious is what this reveals about “ought” judgments, namely that the two divergent perspectives she entertains diverge on the preservation of OIC. Either participant 36 is being asked about logistical possibility—in which case disagreement would be her response—*or* some ethical consideration—in which case agreement would be her response—both of which preserve OIC. For this response to be an interesting challenge to OIC, then participant 36 would need to agree with the statement as an ethical claim *in light of* the logistical impossibility; but that isn’t the case.

Lastly, some participants seemed to be basing their “ought” judgments to some degree on whether Brown wants or intends to make the meeting. We might not typically think of obligations as depending on desires or intentions, especially when obligations often go *against* our wants: *you*

don't want to go, but you ought to since you promised. Participant 28, for example, produced this exchange during her interview:

Me: And what did you respond for question 3 and why?

Participant: Neither agree or disagree....ummm....the way I interpreted that question was he ought to make it that that means like he he can make it but if he wants to in my head that's how it work so I neither agreed or disagreed because it was really up to him until the end....ummm..and s...I feel like subconsciously he coulda made the decision to either go or not go.

Me: So it was based on whether or not he wants to go [participant says "Mm-hmm"]...had you known that Brown didn't want to go and that's why he passed the time would that change your response?

Participant: Yes because I woulda said he shouldn't have gone I feel like.

Me: For question 1 you said that Brown you agreed that Brown can't make it to the theater by 6 [participant says "Mm-hmm"]...umm...and then in question 3 you mentioned something about he's able to...what's the relation between your response to question 1 and question 3?

Participant: Well question 3 he's able to but then he decides that he doesn't want to so therefore he can't make it to the theater by 6 for question 1.

Me: So for question 3 you said he is able to so is there a certain...when when are you thinking question 3 is asking about or is it about a time or not about a time?

Participant: I feel like question 3 it's at I feel like question 3 is applied at like 5:45 because at fo..5:45 he can decide to leave or he can decide to stay longer and then question 1 is specifically at 5:50 so.

Me: Perfect so you were th...you were answering question 3 with that moment of choice in mind [participant interrupts with "Yes"] at 5:45 [participant says "Yes"] to go to not go he decides to wait around [participant says "Mm-hmm"]...perfect so then as a follow up would you agree or disagree with the following statement at 5:55, Brown ought to make it to the theater by 6.

Participant: At 5:55 yes...he ought to make it to the theater by 6 5:55? [I say "Yea"] Well no he can't make it 'cause it takes him 15 minutes to get there right yea...⁴⁵⁸

My focus here is on those moments when participant 28 connects "ought" to "want" as if to say Brown "ought" to go just in case he wants to. This understanding of "ought," while perhaps counterintuitive in one sense since obligations often go against wants, is meaningful in the sense that people ought to do the things they want. That is, it seems to treat "ought" in part as a measure of what one determines for herself. Of course, it is difficult to get the specific contours of how this participant is using "ought," but it reads to me as the judgment that if Brown doesn't want to go, then it's not clear he ought to.

In sum, there are just too many interesting ways people understand “ought” within the confines of superficially simple scenarios like Brown’s for researchers to view a quantitative endorsement of obligation to be viewed in isolation as anything but opaque. This suggests that future x-phi research on OIC needs to carefully account for the many ways laypeople understand “ought” in connection to blame and ability. Interestingly, Kurthy and Lawford-Smith, as previously mentioned, feared something similar regarding the ambiguity of “ought” in relation to a different study on OIC (“Ought, Can, and Presupposition: An Experimental Study” by Mizrahi). Regarding a set of vignettes, in which someone either forgets about a meeting or finds themselves locked in a room and unable to attend, Kurthy and Lawford-Smith worry about the possible definitional interpretations: 1) that participants might think of a question of ought as a question to whether an agent “has a reason” to meet an obligation⁴⁵⁹; 2) that participants with “rule-based moral systems” might view questions of “ought” as asking about whether there is—or I might add, ought to be—“a rule in favour of keeping appointments, or keeping one’s word”⁴⁶⁰; 3) that participants might view ought “as backward-looking rather than forward-looking” such that an agent is blameworthy for not being able to make a planned meeting.⁴⁶¹ While I didn’t test the vignettes that these researchers speak to, I do feel that I vindicate their worries in looking at those from Chituc et al., especially because I found something similar to some of these very definitional understandings in the vignettes I did explore. And, as explored above, I found *even other* interpretations of “ought.” For just one example of a similar definition, recall that participant 33 selects 5 to Brown ought to make his meeting for the following reason: “I figured the meaning of the question is...uhh... Brown has a responsibility to get there by 6 or Brown.”⁴⁶² This is quite similar, if not identical, to the notion that one uses “ought” to express that there is or ought to be a kind of rule obtaining for those who make promises such that they gain a responsibility in such

situations. Or at the very least, this response, and the surrounding interview transcript (available in the previous suggestion section from this chapter), displays a kind of thinking very much in line with rule-based morality in that those who make promises are responsible to uphold them. Later in the interview, participant 33 says as much: “yea based on the fact that he promised and that he could’ve planned and should’ve planned things in advance I concluded that he ought to have been there by 6.”⁴⁶³ My data in this section, then, validates Kurthy and Lawford-Smith’s general worry about ambiguity. If we can’t verify how participants understand “ought” then it is easy to confuse this or that quantitative response with this or that specific use of the term.

FIVE; DIRECT V. INDIRECT CONTROL OVER ACTIONS MATTERS FOR OIC

Here I suggest that quantitative responses corresponding to concepts of “ought,” “blame,” and “can” might be opaque to whether or not a given participant is judging an agent’s action based on direct v. indirect control over certain events such that two identical quantitative surveys might each reflect divergent judgments. This observation links up with research on direct v. indirect responsibility. Wesley Buckwalter, in his paper exploring implicit attitudes and moral responsibility, outlines the distinction between these two kinds of responsibility as stemming from the difference between “the kind of control we have over the behavior in question.”⁴⁶⁴ For example, we might not expect a professor to have direct control over her implicit biases toward specific students when grading papers, but she exerts indirect control when she opts to anonymize her students’ work prior to grading.⁴⁶⁵ While Buckwalter’s project is different than mine, his focus on this point is useful to my suggestion here that direct v. indirect control matters for how participants understand OIC. Before getting to my qualitative support for this distinction, it is helpful to see how Buckwalter deals with the topic, albeit for different purposes.

Buckwalter looks to empirical research on OIC as it is relevant to his discussion on implicit attitudes and moral responsibility, a project distinct from mine. Nonetheless, he points to the OIC study, and OIC experiment 3 in particular, as one of multiple studies that demonstrates the intuitiveness of the claim that in some cases agents can be morally responsible for actions they can't carry out.⁴⁶⁶ Buckwalter then highlights the important distinction between direct and indirect responsibility as it connects with the ability condition, or the condition that agents must always be able to carry out actions they are morally responsible for.⁴⁶⁷ To do so, he specifically references OIC experiment 3 and a similar case from Walter Sinnott-Armstrong⁴⁶⁸ in which Adams bails on his promise to Brown by catching a movie across town; Buckwalter notes that the findings from Chituc et al. put pressure on the intuitiveness of the ability condition as called into question by the case of Adams bailing on Brown.⁴⁶⁹ He also, interestingly, offers an important potential objection to the conclusion that the cases with Adams and Brown are so threatening to OIC:

It might be objected that while protagonists like Adam [*sic*] and Brown do have moral responsibilities in the cases [referenced], they did have at least some ability to fulfill them. For example, though Adams lacked the ability to fulfill his promise at the end of the story, he did have the ability to control this earlier in the story. Recalling the definition above, if it is reasonable to expect Adams to control his behavior at this earlier point, then Adams has “indirect” responsibility for his behavior at the later point.⁴⁷⁰

This bit is now moving toward what I am interested in with respect to Buckwalter's project. Buckwalter posits that if this appeal to indirect control holds up, then “it suggests that the ability condition is false because agents can be morally responsible for things beyond their direct control in virtue of being indirectly responsible for them.”⁴⁷¹ In simpler terms, Buckwalter writes, “this suggests that moral responsibility is compatible with inability when the source of an inability can

be traced back to the agent in certain respects, such as their free choices or negligence.”⁴⁷² Immediately following this, Buckwalter suggests that this appeal to the difference between direct and indirect control could be invoked to defend OIC—certainly an interesting proposal—though he wonders its efficacy in backing “the argument against moral responsibility for implicitly biased behavior”.⁴⁷³ And here’s where Buckwalter’s work most clearly connects to mine: Though my goal isn’t to defend OIC, I am arguing that Chituc et al.’s findings don’t obviously negate OIC as they seem to think, and there’s reason to think direct and indirect control is relevant to the discussion. So, I will no longer focus on Buckwalter’s (different) project *and I will* take him up on this implicit suggestion that I might defend OIC in light of this distinction between kinds of control and responsibility. More specifically, I will examine how these kinds of control might be influencing participants who encounter cases involving Adams, Brown, and a missed appointment.

It occurs to me that Buckwalter has just proposed what might be the best explanation of the data from Chituc et al.: what Chituc et al. take to undermine OIC is really just a more nuanced understanding of “ability” such that ability isn’t only understood in some specific moment but is sometimes attributed to a series of causal events such that being able to do something at 5:25 to prevent lateness at 6:00 is viewed as a kind of “ability.” Of course, Chituc et al. might object and claim that, nonetheless, participants who select that Brown can’t and ought make it in OIC experiment 3 are attributing moral obligation *at a moment* when ability is now long gone. But this assumes that what people really mean to express is that at the precise moment that agent is unable to carry out a promise that what they ought to do is carry it out. What might plausibly be going on is that some people hold a more nuanced view wherein agent *had* the ability to prevent a failure of obligation and that despite the fact that they no longer have such ability, they bear an indirect responsibility for either 1) not anticipating the external events that caused their inability to occur

or 2) intentionally doing those things that would bring about that inability. And that while they view Brown as holding that indirect responsibility, and agree that “Brown ought to make it to the theater by 6,”⁴⁷⁴ said participants have difficulty saying that agent still “ought” to make it on time when asked about a concrete time after the inability is in place when there is no real window for changing behaviors, such as 5:55. Quantitatively, this response would look like the following: someone on the high fault version, wherein there is more desire to hold Brown indirectly responsible, chooses that Brown can’t at 5:50, is to blame, and ought to make it to the theater by 6. Qualitatively, this response would be that same someone back-tracking on the ascription of “ought” in the face of the timestamp, 5:55. And, we have just that in participant 30 whose full transcript is presented and analyzed in chapter six. It’s true that participant 30 doesn’t directly state that Brown has “indirect” control, but this understanding tracks her responses in that she attributes blame to Brown just for consciously bringing forth his inability to make the promise:

Me: Excellent. Do you agree or disagree with the following statement? At 5:40, Brown is to blame for not making it to the theater by 6.

Participant: Mmmm...[pause]...I..I guess I’m thinking that it would depend on when he decided he didn’t [laughs slightly while speaking] want to go....ummm...like if he decided at like before that time that he didn’t want to go then yes if he decided at 5:45 then I would say at 5:40 he isn’t to blame for not making it ‘cause he was still technically coulda been on time.⁴⁷⁵

That is, both of the following are the case: 1) Participant 30 says that Brown’s decision to not go is what makes him blameworthy; 2) This decision is the thing over which Brown has indirect responsibility at a later time.

Another way of seeing evidence for participant concern for indirect responsibility is catching a participant’s thinking prior to his or her reading the follow-up questions, and thus prior to his or her thoughts being funneled through the follow-up questions. Judgments about direct and indirect responsibility are nuanced, but no such nuanced response is available in the follow-up questions—i.e., someone can’t choose *Brown no longer can make it on time, and so in that moment*

he ought not do what he can't, but the bigger picture is that he ought not to have waited around too long in the first place and he is responsible for that. So, it appears that some judgments are problematically funneled through three limited questions and these responses are being interpreted incorrectly. How can someone express, on these questions, that Brown ought to have taken more precautions in the past to prevent certain pitfalls? If that is what a participant wants to express, which seems natural, then they are forced to funnel that evaluation on the situation through the limited questions of agreement and disagreement on three concepts which don't really leave room for commenting on Brown's more direct and indirect control over previous actions. For example, participant 7 gives a brief reaction to Brown's situation during the think aloud session right before reading the survey questions and therefore right before his reaction gets funneled through the questions; this reaction appears to comment on Brown's direct control over actions prior to his inability to make the meeting (which means he has indirect responsibility for his later inability to make the meeting): "He passes the time for 5 minutes, so that he will be unable to make it to the cinema on time. Because Brown decided to wait, Brown can't meet his friend Adams at the movie by 6. Why didn't Brown just call Adams and say hey I don't want to see the movie...have fun without me..."⁴⁷⁶ And participant 3 let's out a similar reaction before reading the questions: "Because Brown decided to wait, Brown can't make...can't meet his friend Adams at the movie by 6. Well, he should tell his friend."⁴⁷⁷ But the three questions on the survey don't really allow for the kind of response that comments on what Brown should have done prior to his inability or even in light of his inability, namely pick up the phone. To have subjects cram their judgments through these sterile OIC questions is limiting to what appears to be one straightforward response: call or text to cancel.

What follows are more moments from the qualitative data that speak to the role indirect responsibility plays in people's judgments. First, participant 8 references the possibility of Brown's indirect control in the low fault case since failing to meet an obligation due to car trouble isn't always a blameless occurrence:

OK, question 2. Brown is to blame for not making it to the theater by 6. So, none OK...so, if he had decided to leave a little bit earlier maybe he could have made it to the theater despite all of these alternative things that end up happening such as it being cold and that dissuading him from riding his bike and his car leaving..ah..breaking down..so if he'd chosen to leave a little bit earlier maybe he could have may..may...made it on time..that's not to say that he's at fault for not leaving earlier...so...his car broke down which is not under his control more than likely unless he decided to defer maintenance or [laughs while speaking] something..umm..and it being cold is not his fault either so I kind of think that it's just bad luck in this case or sort of a..he's kind of a victim of circumstance..umm..so..I wouldn't say that he's completely without blame because there were some maybe steps he could have taken to ensure that he was there by 6 but I also wouldn't blame him for it so I am going to...umm...let's see...I still think that I want to disagree but it's gonna be..umm...sort of a a little bit lighter of a disagree so I think I'm gonna put a negative 3 because there were some steps he could have taken perhaps to ensure that he got there.⁴⁷⁸

In addition, we see participants thinking about indirect control in that Brown leaves with just enough time to arrive. This, by the way, exposes some absurdity of this thought experiment in that Brown has literally 15.000 minutes to make it with a hard cutoff on the theater letting people in and he leaves with exactly 15.000 minutes ahead of him such that the slightest inconvenience would make him late. What kind of whacky movie theater is this? Participant 7 succinctly puts it,

“The cinema has a strict policy of not admitting anyone after the movie starts and the movie always starts right on time...where the hell is he going to see his movies?”⁴⁷⁹ And what kind of person leaves with literally zero seconds to spare? If we take this scenario at face value it sounds like Brown has his ticket ripped by the usher at the last moment in time the usher would have been willing to do so; had a slight oncoming wind slowed his car but a nanosecond, he would have encountered a stern, finger-wagging usher rather than the movie he already *paid* to see. Here are a few participants worried about this lack of margin for error; first participant 5 who relates to this kind of punctuation:

Me: To begin, please tell me about your first reaction to the scenario you read about with Adams and Brown.

Participant: Umm..the first thing I thought about was how he left 15...exactly 15 minutes earlier...umm like with it the exact amount of time he would need given every situat... or every scenario like going perfectly umm..it made me think about that's exactly how I de...how I live like I always give the perfect amount of time and...a lot of the time that's not enough time so if he left a little bit earlier umm...he may have prevented the situation umm..and I also thought why why didn't er maybe he did but why wasn't it in there he contacted his friend that he wasn't gonna make it.⁴⁸⁰

And another example from participant 15 when asked about question 2 (on blame):

Me: And what did you choose and why?

Participant: Umm...I chose 0 neither agree nor disagree because...umm...he didn't leave early so I guess it's his fault for not preparing for any disaster but he also left on time so I think it's kind of fair...umm...yea.

Me: Initially you had put negative 2 and [participant interrupts with “Mm-hmm”] then you went back and changed it what explains your re-thinking of the question?

Participant: Right, well I re-read through it and...ummm..sort of like thought more about how much time it took him to drive to the cinema versus what time he actually left and he left at 5:45 and it takes him 15 minutes to get to the cinema so he should get there exac...exactly at 6 so they ideally would be able to let him in but that's not leaving himself any time like even a minute so it was a little bit his [laughs slightly while speaking] fault.⁴⁸¹

In each case, indirect control operates in the participant's reasoning. To strip that away because their quantitative responses don't reflect that—and how could they?—seems problematically reductionistic.

Having now discussed both temporal concreteness and indirect control, I might add in that Hannon's study seems to corroborate one or both notions; I'm not yet decided on how to think

about it. That is, Hannon's findings suggest that some people mean to express that agents *ought to have* fulfilled a promise when they agree that the agent ought to do so.⁴⁸² This comments both on indirect control, as it seems to imply agents ought to have done those things they previously had control over in order to meet the obligation, *and* on temporal concreteness because language about what someone ought to have done often uncouples the ascription of obligation to the same moment of in time as the inability (thus preserving OIC). In sum, Hannon's experiment further iterates how these factors of indirect control and concreteness seem to influence people's judgments related to OIC.

SIX; MORAL REMAINDERS MIGHT EXPLAIN SOME OIC JUDGMENTS

I also suggest that participants' quantitative responses might be influenced by moral remainders such that what appears to be a straightforward ascription of obligation to an agent might really be an attempt to ascribe a moral remainder based on the limited options available on a survey. Moral remainders, or moral residue, are discussed within the philosophy on moral dilemmas, but they can be understood in other contexts. Rosalind Hursthouse helpfully defines the idea in her work on their role in exploring moral dilemmas from deontological, utilitarian, and virtue ethics positions:

Suppose someone is faced with a strict moral dilemma in which x and y are equally bad. Then, whichever they choose, they choose what is *ex hypothesi* bad, and we expect them (especially when we think in terms of real examples) to register this fact (or thought) in some way—by feeling distress or regret or remorse or guilt and/or (in some cases) by recognising that some apology or restitution or compensation is called for. This—the

remorse, or regret, or the new requirement to apologise or whatever—is called the (moral) remainder or residue.⁴⁸³

Of course, it is debatable that Brown in OIC experiment 3 ever finds himself in a dilemma: neither being stuck at home with no time to make it to the theater by 6 nor being stuck on the side of road seems to present him with any real choice as to whether to uphold the specific promise he made. And yet, this notion of moral remainders straightforwardly applies to Brown because it captures what some people expect agents to experience when a promise is on its way to being broken, whether from intentional or external forces.

When the car breaks down, some of us expect an agent to feel obligated do the next best thing, which is one kind of remainder, and to feel a kind of psychological guilt, at least temporarily, for standing up a friend. (This first kind of remainder, the next best thing obligation, again looks like Kurthy et al.’s discussion of the “scope-shifting problem” wherein participants see obligation shifting from carrying out the original obligation to finding some good alternative in light of the onset of inability.⁴⁸⁴) In fact, some of us would find it a bit troubling if a friend whose car broke down didn’t feel any psychological hold from the obligation she made. Of course, we’d not want her to feel *that* bad if it is just movie plans being thwarted, and instead we’d hope she worries about her safety more so. Nonetheless, to learn that the friend felt perfectly relinquished from considering all alternative actions when the car breaks down, such as calling and alerting us about the incident, and to learn they felt totally released from any psychological sense of guilt or apologetic-ness would be, in many cases, almost worrisome for the friendship. Does this friend really feel so little commitment that she pays not one thought to my being stood up at the theater? And when this friend actively chooses to wait too long so as to miss the movie, we’d hope she feels something more. She should feel bad for intentionally bringing about the broken promise, or

at least she should feel some kind of duty to repair the situation, or make up for it in the future. Again, to learn that she felt or thought nothing is almost as troubling as her original decision to break the promise!

If a participant in OIC experiment 3 is trying to convey that a moral remainder applies or a residue ought to exist for Brown given his actions, then that participant might agree that Brown ought to make it in question 3 in order to imperfectly convey that a remainder remains. To a degree, Chituc et al. grapple with a worry like this when they reference *blame validation*, which they define as “a process in which a motivation to blame can increase judgments of ability and responsibility—suggesting that when someone is blameworthy, participants may respond by exaggerating their obligations to hold them accountable,”⁴⁸⁵ and *excuse validation*, which they define as the phenomenon “where a motivation to withhold blame leads participants to deny that a rule has been broken—suggesting that when someone is blameless for a transgression, participants may respond by downplaying their obligations to protect them from censure.”⁴⁸⁶ Specifically, they note that objectors to their work might claim that participants had “distorted judgments” as a result of blame validation, but they counter by noting that the same correlation between judgments of “blame” and “ought” appeared in the low fault cases where we wouldn’t expect blame validation to distort those judgments.⁴⁸⁷ A different explanation, they suggest, is that philosophers who support OIC might be biased by excuse validation such that they negate an agent’s obligation in occasions where an agent lacks ability due to no fault of his own so as to protect him from blame or punishment.⁴⁸⁸ This is certainly a plausible account. But it assumes that participants’ ascriptions of obligation are reflective of the kind of obligation to do just the thing that agent promised to do and not that they feel that the residue or remainder should be more or less strong or present for the promise-breaker. In other words, we might better explain the

relationship between blame and ought judgments by hypothesizing that people want more remainder—more guilt, reparations, apology, alternative action—in cases of high fault than they do low fault such that they are taking the opportunity to express the need for this residue in agreeing more or less with the question of “ought” since the question is vague enough to be interpreted as such an opportunity.

Chituc et al. might object and argue that what I see as participants hoping to ascribe moral remainder is no different than the blame validation objection they take themselves to dispel. But whereas with the blame validation theory we might expect the low fault versions to produce no such observed relationship between ought and blame, that isn't the case for moral residue. Even if someone isn't to blame, we can still want them to experience or be subject to a residue. Even if your car broke down right after you got it fixed, I still expect you to feel or think with some kind of residue in the form of guilt or of brainstorming new ways to get me to the airport on time—you made a promise, I relied on you, and I'm about to miss a flight. (Of course, I would personally feel a strong need to reassure you that you needn't feel such guilt; after all, I knew there was a chance that in relying on you things wouldn't pan out and you did take reasonable precautions to ensure effective transportation.) So, whereas with low fault we should expect blame to taper off, we shouldn't necessarily for residue ascriptions. Some participants might have upheld what appears to be a relationship between blame and obligation, but was really a relationship between blame and residue-ness.

This hypothesis of remainder being masked as obligation is something that could be explored in future research. For now, though, I have some qualitative data that speaks to it. There is evidence from participant 8, who completed the low fault survey with 5 (Brown can't), -3 (Brown is to blame), and 3 (Brown ought). Here is the first excerpt of interest:

Me: Great. And what did question 3 on the survey ask?

Participant: Question 3. Brown ought to make it to the theater by 6.

Me: And again, what did you choose and why?

Participant: I chose plus 3..umm..I was a little thrown off by the word like ought like he should make it there by 6...umm...and so I think that because he was super excited about the movie and because he had committed to seeing this movie with his friend..umm..because he wanted to do this then all of that leads to he should have made it there by 6 such that he could see the movie..umm..but at the same time it's hard to say that anyone should be kept to a certain commitment or a certain..uh..whatever this could be considered sort of you know date..umm..plan because some things will happen and other types of incidents will get in the way and so should is just a hard or ought is just a hard word to work with in terms of when things other things happen what can you do so that's why I put a plus 3 because I sort of wanted there to be a little bit room of room for other things to happen.

Me: So in line with that last question..[participant says "Mm-hmm"]..do you agree or disagree with the following statement? At 5:55, Brown ought to make it to the theater by 6.

Participant: So in line with that question at 5:55 he ought to make it to the theater by 6. [pause] Well, hmm..[laughs] so at that time he wouldn't physically be able to make it to the theater by 6 because of his transportation limitations but the should have feeling the ought to have feeling is still there..umm..and so I would say no he shouldn't have to make it there by 6 because he physically can't..umm..but I think in terms of that feeling of commitment that would still be there at 5:55.

Me: Great, so your response on the survey for question 3 was that you agreed that Brown ought to make it there [participant says "Mm-hmm"] by 6 and then the question I just asked about at 5:55 you said you disagreed [participant says "Mm"] would you mind elaborating a little bit more on that?

Participant: Yea yea because at 5:55 it's gonna be physically impossible for him to make it to the theater by 6 so in terms of asking should he get there by 6 as in like is that expected for him to get there at 6 right is this possible that would be a no but in terms of the feeling as though he should have made it there in time feeling as though he's still committed to this..umm..that would be that would be a yes so it's kind of like a is it physically possible versus is this still something that I feel committed to and like I should have been able to make it happen..um...so I guess it's just interpretation of wording but there is that distinction between physical possibility and then what your brain is thinking about what could have been and you should have done to change that.⁴⁸⁹

First, it's noteworthy that while participant 8 selected survey responses that appear to negate OIC—i.e., Brown can't and ought to make it—she describes her reasoning with more nuance than this captures. She affirms that Brown has an obligation of a kind, but she leaves room for error given that things come up that prevent us from fulfilling our obligations. Already we see that she might actually be *supportive* of OIC despite her quantitative responses since presumably the endorsement of "ought" was lessened because things coming up equates to things preventing someone like Brown from carrying out his obligation. From here, I ask participant 8 about the more concrete moment of 5:55 and things get more interesting and noteworthy. In response to this temporal concreteness, participant 8 makes a distinction between what Brown psychologically

experiences, a sense of obligation (residue), and what he should or shouldn't do. In the end, she denies that Brown ought to make it precisely because he can't, but that psychologically he will still feel obligated at that time (residue). Before moving to the next relevant excerpt, which picks up on this interesting distinction, I want to note that participant 8 highlights in this something like a meaningful definitional distinction in that she points to two plausible interpretations of Brown's obligation as described by "ought": one physically bound and the other psychologically experienced. While she doesn't specifically tie her response to two different definitions of "ought," the surrounding context implies that "ought" is the concept from which these definitions stem.

Later on in the interview, this interesting split between the psychological and the physically constraining is revisited:

Me: OK..and..so y..for question 1, uh, you agreed that at 5:50 Brown can't make it to the theater by 6 [participant says "Mm-hmm"]. And for question 3 you agreed that Brown ought to make it to the theater by 6. Looking at these two responses together, do you agree or disagree that Brown ought to do something that he can't do?

Participant: [whispers "ought to do something that he can't do"]..uhh...that's OK so I would say that I agree that if it is physically impossible for him to do something then I understand that he ought not do it physically..right..that doesn't work..umm..but I do think that he psychologically would still feel as though this is something that he should have been able to do should do..um..and so there's sort of a difference between what he can physically do and then how he thinks..um..that what he thinks that he should be able to do...um..but but sort of I suppose as an overall thing if he can't physically do it he shouldn't do it even if he feels committed.⁴⁹⁰

Here I ask participant 8 to follow up on this interesting distinction in her responses and she affirms once again that physically Brown can't and so he shouldn't, but this time she adds that this is the case *even if* that psychological experience of obligation persists. Participant 8's quantitative responses merely appear to negate OIC, but when a single important concrete detail is added—i.e., a timestamp that actually pits "ought" and "can" against each other—she affirms OIC in saying, in so many words, that Brown shouldn't have to make it. At the same time, she introduces a possible explanation why someone might judge that Brown ought to make the meeting: regardless

of whether or not he is morally required to make the appointment, perhaps we expect him to *psychologically feel* the obligation (residue).

What is also interesting with these responses is that participant 8 dealt with low fault version. That is, where there might be no great reason to *blame* Brown—which participant 8 affirms with a -3 on blame—there is still a moral residue or remainder to be attributed to Brown. Again, this suggests that this participant, and perhaps others, are happy to say Brown “ought” to make it by 6 because they want Brown to *feel* responsible even if from some ethical reasoning standpoint he no longer is. But, when asked about 5:55, a concrete time which seems to ask if Brown should literally do what he can’t do, the interpretation of “ought” changes to one concerned with remainder as well as one concerned with ethical reasoning as such. Interestingly, participant 8 still attributed some blame, or not full exoneration, to Brown despite the fact that she reacted to the low fault scenario: “I wouldn’t say that he’s completely without blame because there were some maybe steps he could have taken to ensure that he was there by 6 but I also wouldn’t blame him for it so I am going to...umm...let’s see...I still think that I want to disagree but it’s gonna be..umm...sort of a a little bit lighter of a disagree so I think I’m gonna put a negative 3 because there were some steps he could have taken perhaps to ensure that he got there.”⁴⁹¹ Still, this doesn’t seem to be the kind of blame sufficient for attributing blame validation, so participant 8’s responses appear best understood in terms of a desire to ascribe something like moral remainder or residue to Brown.

What this all suggests is that some participants, such as 8, might be agreeing that Brown ought to keep his promise because they agree that he ought experience a residue or that he ought be under obligation to make reparations (remainder). While not all participants have been so

explicit as participant 8, the effects of remainders nonetheless could be operating in the minds of more participants, bearing on their quantitative selections.

SEVEN; WE NEED TRIANGULATION TO CAPTURE THE COMPLEXITY OF OIC

In total, I've suggested that quantitative surveys are tools too blunt to capture the complexity of people's judgments of obligation, blame, and ability. I have provided evidence that people have complex understandings of these concepts that keep rearing up when participants are asked how they think about complex scenarios in which obligation and ability are put in tension. To contribute to the OIC debate in the literature, my positive claim is something like this: people appear to have rich and complex judgments relating to obligation and ability that we've only begun to fully grasp. Specifically, we saw that individuals express complex ought/can judgments and that there is interesting variation between individuals' judgments. The conversation in x-phi of OIC might benefit from appreciating some inherent tension in the problem of OIC that isn't resolved by asserting whether or not OIC is or isn't correct, or is or isn't intuitive.

The solution that I propose to my suggestions is to take them seriously in the form of more empirical and analytical investigation, in survey and interview formats, with a focus on triangulation of methods.⁴⁹² Perhaps the next projects in the x-phi of OIC might develop surveys based on my suggestions to see if certain of them can be borne out in quantitative questions. And from there, further think aloud sessions and interviews are in order. But to treat a handful of surveys that have face validity to philosophers as revelatory of people's complex understandings of moral obligation is problematic. It ignores the methodological problem of experimental philosophy by assuming that what we think we get in a survey is really what we get. As we have seen, the ideal story is just that, ideal.

CHAPTER EIGHT:

Conclusion:

Ethics is Messy, Triangulation is Worth the Trouble

I. The Messiness of Ethics and the Limitations of My Work

This project was aimed at learning more about how people think about issues in ethics, and bringing those findings into better conversation with interested philosophers. And a big part of that undertaking was not only learning about “ought implies can” and how laypeople conceive of it, but also in analyzing the tools that best help experimental philosophers in their inquiries related to moral philosophy. In some sense, I’ve dealt with a very messy problem—i.e., seeking understanding of philosophical thought in laypeople—that stems from a messy object of inquiry—i.e., the rich and nuanced ways people think about issues in ethics—to ultimately prescribe a messy solution—i.e., using multiple methods to triangulate on said thought processes. And so, it shouldn’t be surprising that there are some significant limitations not only to the work I analyzed but to my own.

My experimental study in particular was driven by a number of objections to (or questions for) the ideal story of quantitative data collection in certain χ -phi studies. However, by introducing new research methodologies, my research is also open to a number of objections, or at least reservations, of its own. First, it is true that bringing participants into a laboratory setting for a think aloud session or interview is time-consuming and expensive. I paid each participant \$25 dollars for his or her time, which is much more than the fractions of a dollar studies often offer for the completion of a survey on Amazon Mechanical Turk. Even if they send out ten times the

surveys I had completed, these qualitative methods are still much more expensive. I felt a need to pay as much because I was asking as much from participants.

There is also the time and work needed to conduct the interviews and transcribe them in addition to the work already baked into collecting and analyzing quantitative survey data. For those who have never transcribed audio before, I assure you this process is exceptionally time-consuming and frustrating. In total, then, the methods I advocate for as part of triangulation are expensive and time-consuming. However, I think they are worth it in exposing where we might make changes to future quantitative surveys and in shedding light on whether or not we are getting the kinds of judgments we think we are in ethics and beyond. As I have argued, prior to using such qualitative methods, experimental philosophers are subject to the MPEP in that they can't justify, based on their limited research methods, that they have data reflective of people's judgments as they claim to have them. I have specifically shown that experimental philosophers of OIC have not overcome the MPEP. It was by conducting this more expensive study that I came to argue this. Using qualitative methods, researchers might save more time and money in the long-run: sending out survey after survey on Amazon Mechanical Turk might be wasteful if those surveys don't reflect the judgments we take them to reflect.

The next core limitations pertain specifically to the think aloud method. Van Someren et al. discuss two “threats to the validity of reports” when employing the think aloud method as opposed to other methods.⁴⁹³ First, they note that “thinking aloud takes place concurrently with the cognitive process,” which ultimately means that the “cognitive process takes longer” and that “people are able to slow down the normal process to synchronize it with verbalization.”⁴⁹⁴ This alone I consider to be a potential problem, as it admits to some interference with a subject's “normal” thought process—unless a subject almost always thinks aloud through her problems in

real life. In line with this, van Someren et al. claim that “subjects frequently report that sometimes verbalization does not keep up with the cognitive process” which renders a subject’s report somewhat or occasionally incomplete.⁴⁹⁵ This strikes me as a limitation without a perfect solution, but my hope was to fill such incomplete parts of the think aloud sessions by asking questions in the subsequent interview I conducted with subjects, thereby mitigating what might have been lost through or altered as a result of thinking aloud. Second, van Someren et al. point to a problem that results from a person working to verbalize his or her thinking which is not tailor-made for verbalization and how such an act interferes with a person’s working memory: “If reasoning takes place in a verbal form then verbalizing the contents of working memory is easy and uses no capacity of working memory. However, if the information is non-verbal and complicated then verbalization will not only cost time but also space in working memory because it becomes a cognitive process by itself.”⁴⁹⁶ That is, trying to verbalize certain thought processes might significantly interfere with the resulting protocol. This is especially problematic for my research ideas since experimental philosophers are interested in a subject’s thinking generally (and often the immediate intuitive thinking more specifically), not solely their readily verbalizable thinking. I think this too has no perfect solution. However, the fact that we have a transcribed version of participants’ protocols should alleviate some worry about interference since we can see if and when people appear to stammer, stumble, or even express difficulty in thinking aloud.

Similarly, there are also limitations inherent to conducting interviews that apply to my study. Try as I might to remain an objective interviewer asking questions about Brown and Adams, I am prone to bias as anyone else is. I certainly might have emphasized questions or highlighted certain responses that align more so with my thinking about matters of obligation, blame, and ability. To account for this, I tried to put large excerpts and full transcripts to show how a

conversation proceeded and leave the reader to judge for herself if I was being more or less biased. At the same time, I hope that other researchers conduct similar interviews so we can see if they get a similar array of interesting responses. In the end, there are enough examples in the transcripts to show that people's thinking was far more complex than would be captured by just the limited survey questions presented by Chituc et al. Even when biases might be influencing an interview, we have the think aloud transcripts which include someone's thinking without my intervention. And there alone we see this complexity on display that suggests the MPEP is a problem for survey methods alone.

The penultimate limitation I want to address, which will take some space to unpack, results from the fact that I didn't sufficiently trust my own intuitions, shaped by my unique upbringing, when conducting my literature review and my study. My mother is from Oklahoma, and I've come to realize over the years that many words and phrases I use with regularity are unique to my mother's culture and dialect. I'm fairly confident that the way I understand "ought" has been influenced by my mom's culture and from my spending time around family from Oklahoma and Kansas. Today, the word hits my ears with at least two meanings, one relating to moral obligation—for this brief discussion we can call this the "moral ought"—and one relating to the logistics of accomplishing something—we can call this the "logistical ought."⁴⁹⁷ The moral ought is what I've been discussing for most of this dissertation: 'You ought to go; you promised you would!' The logistical ought strikes me as non-moral, or at least not necessarily moral, and conveys that the speaker, as far as he or she can tell and has calculated to be the case, thinks something will occur as predicted or that something is sufficient for accomplishing some end: 'If we leave now, we ought to make it on time'; or 'One gallon ought to be enough to get you there.'⁴⁹⁸ It conveys that, *logistically* speaking, something will happen or will suffice. And while these two

meanings were occasionally on my mind in designing my study and reading the OIC literature, I didn't sufficiently account for the logistical ought when analyzing other OIC studies or when crafting my own. While I made significant efforts to account for the fact that "ought" can have multiple meanings in different contexts—as is evidenced by my analysis throughout this dissertation—I didn't focus enough on this logistical meaning, and I didn't realize that until my project was largely complete. Perhaps this is because I doubted that other people were used to hearing "ought" this way, or because I didn't want to bias participants in favor of using "ought" in any particular fashion; I honestly don't know. But for whatever reason, I didn't take seriously enough my own intuition that "ought" might take on this meaning for some participants. For the clearest example of this limitation on my work, consider my interview question to participants as to whether Brown ought to make his appointment at 5:55; for those who understand "ought" in both the moral and logistical ways, this question is potentially opaque or confusing. Therefore, I ought to have better accounted for this reality, and now I see it as a limitation to my work that I want to account for in future projects on OIC. As a starting point, I would want to learn more about this logistical interpretation of "ought" by researching whether it is more often connected to a Southern dialect or if it has more widespread use; I can imagine a follow-up question on a survey that asks participants about their familiarity with specific definitions of "ought" such as the logistical one.

Now, this limitation doesn't largely affect my criticism of the work by Chituc et al. given that one of my worries about their experiments has been that people seem to understand "ought" in multiple ways. However, I consider my lack of focus on this logistical meaning as a limitation to my own work in that I could have inserted interview questions that would have better teased apart which "ought" someone was thinking about when discussing Brown at 5:55. For example, I

could have swapped “ought” for “obligated” in reference to Brown at 5:55 in some interviews to see if I would get similar responses.⁴⁹⁹ While Chituc et al. didn’t vary the language of “ought” in their study, some studies in the OIC literature did. For example, Buckwalter and Turri administer survey questions in their different experiments that ask if agent “ought” to do something, if agent “is obligated” to do something, and if agent “has a duty” to do something.⁵⁰⁰ So, in focusing on the Chituc et al. study, I ended up working solely with “ought,” but I think future x-phi work on OIC, whether conducted by me or other experimental philosophers, should both vary verb (or verb phrase) choice for “ought” and more fully account for the logistical understanding of “ought” in relation to the moral understanding of “ought”—if memory serves me right, the OIC literature I analyze in this dissertation doesn’t put much focus on the logistical ought. Of course, other issues arise when “ought” is swapped for other verbs (or verb phrases), because it might turn out that people view the moral force or requirement of “ought” to be different than the force or requirement of “obligation” or “duty” even if “ought” is sometimes defined in terms of obligation or duty. Nonetheless, future work on OIC would benefit from varying language in addition to applying qualitative methodologies like the think aloud method and interview.

With all that said about this limitation regarding multiple meanings of “ought,” I do want to emphasize that my appreciation for multiple meanings of “ought” certainly isn’t absent from this project, nor is it the case that my analyses of my qualitative data don’t speak to the possibility of people interpreting “ought” in different fashions. Specifically, my fourth suggestion for future OIC research in chapter seven is that definitional distinctions matter for how people understand “ought implies can.” In fact, participant 36, whose transcript I discuss during that section of chapter seven, seems to contemplate the logistical ought in addition to the moral ought (this is the

same excerpt presented in chapter seven, and the one that gave me the idea to refer to the “logistical ought” as such):

Brown ought to make it to the theater by 6. I’m not 100% sure...ummm..on the meaning of the statement if the statement means Brown should make it in time logistically I would say no I disagree and if it means Brown should [laughs while speaking] make it there from like a [laughs while speaking] an ethical perspective because he promised to meet his friend there...ummm...I would say I completely agree with that so I’m not sure which..ch..Brown ought to make it to the theater by 6 out of respect for his friend completely agree. Is that what the statement was intended to..logistically no he won’t make it there by 6 if he left at 5:50.⁵⁰¹

In chapter seven, I discuss this excerpt as pointing to the fact that there is definitional ambiguity in participant 36’s response, but here, in the context of discussing limitations of my study, I want to add a sharper focus on the language of the logistical in contrast to the moral ought. Notice that participant 36 wonders if she is being asked about the logistical ought—in which case her response is disagreement that Brown ought to make it—or the moral ought—in which case her response is agreement. The fact that her response would change on behalf of a given interpretation of “ought” shows why it is so important to see how participants understand the definition of “ought” and it speaks to the value of looking more closely at the specific distinction between the logistical ought and moral ought. Fortunately, because I made use of the think aloud method, we know which interpretation she ultimately goes with: “OK, I think I will go with...I think I will go with my..uhh....interpretation of Brown should make it there ‘cause he promised to meet his friend there and I would say I completely agree with that he should have made the effort to do that.”⁵⁰²

Without this transcript, her response to the survey question would have been opaque: how did she interpret “ought”?

In discussing the logistical versus the moral ought, I want to mention just one more transcript that I’ve analyzed in this project so far, that of participant 30 (which is available in the first section of chapter six). My goal isn’t to negate the analysis I performed earlier on this transcript, which I hold to, but instead to show that while I didn’t fully account for the particular distinction between the logistical ought and moral ought when planning my interview questions, we can still see that I seem to have successfully drawn out the distinction. During her think aloud session, participant 30 wonders about what “ought” means in the context of Brown’s situation, ends up picking a meaning most closely connected to the moral ought, and then provides this interesting response during her interview when asked if Brown, at 5:55, ought to make the meeting:

Mmmm....I would disagree with that then ‘cause I I I would think that’s that’s kind of meaning that he’s already left or like he’s he’s about to leave and he should make it there in 5 minutes versus like it says it’ll take 15 so I feel like ah..that I would that wording I would interpret slightly differently than how this is phrased ‘cause this is phrased a little bit more generally of like sh..should he like [laughs while speaking] morally get there at 6 or should he like physically like the amount of distance get there by 6 so [unintelligible].⁵⁰³

We might plausibly conclude that the introduction of the 5:55 timestamp causes participant 30 to interpret “ought” as being logistical whereas the general question as to whether Brown ought to make it gives way to the moral interpretation—the surrounding context of this exchange available in chapter six supports this, too. In hindsight, I might have varied the language regarding the 5:55 question to clarify that this is what was going on; still, from the transcript, it seems clear that participant 30 thinks differently of “ought” when asked about Brown at 5:55 than when asked

about Brown without mention of a specific time. One simple explanation for this switch in responses is that any mention of time brings to mind for some people the logistical meaning of ought. Or, it could be that the mention of 5:55, a time when Brown certainly can't make his meeting, makes the logistical ought the only rational or intelligible ought because the OIC principle renders the question of moral ought at 5:55 unintelligible or irrational. Whatever the case, there is certainly more worth exploring with regard to the logistical ought versus the moral ought in the x-phi of OIC. I feel confident that I dealt with definitional distinctions and temporal concreteness in many ways in this project that account for something like the logistical versus the moral ought, but I could have done more in my research design to facilitate exploration of this particular way of teasing apart definitions of "ought." For now, I'll conclude this discussion of this penultimate limitation and advocate for more examination of this issue in the future.

OK, and lastly is the limitation I created in asking participants if they agree that Brown *can't* rather than if they agree that Brown *can*. While I accounted for this in flipping the scale I used to match the one Chituc et al. used, it is still a mistake that limits the purity of my replication of their work.

II. Conclusion

Ethics matters for the whole of humanity; it is an ongoing conversation that is as much the prize (or burden) of specialized philosophers as it is that of farmers, technicians, mothers, fathers, drivers, actors, etc. The value that experimental philosophy offers to ethics, then, is clear: it provides a space and a toolbox for both learning how laypeople think about issues in ethics *and* incorporating their judgments and deliberations into the sometimes echoing chambers of academia, stirring things up in some cases and emboldening foundational ideas in others. I sought to

contribute to the x-phi project not just in capturing and analyzing more judgments on ethical issues, “ought implies can” in particular, but also in expanding the toolbox so that it may better serve the larger conversation that is moral philosophy. And off I went.

My motivation came as much from an interest in people’s messy judgments on issues in ethics as it did from a skepticism toward current x-phi practice. Experimental philosophers often conduct research by presenting experimental subjects with thought experiments in the hopes of causing the subjects to form judgments/intuitions about relevant philosophical ideas. These judgments/intuitions will ideally be that which prompts said subjects to select particular responses on a quantitative survey. In sum, the ideal story goes that a subject reads a vignette and a follow-up question (or more), reacts to the philosophical content within them, forms an unadulterated judgment/intuition somewhere along the way, and then selects a survey response that accurately captures her judgment/intuition without interfering with or limiting her judgment/intuition. However, there are many places during this ideal story in which things might break down. Survey questions might limit a subject’s thinking, a subject might react to unimportant features of the vignette or follow-up question, specific terms or phrases within a vignette or follow-up question might take on different meanings for different subjects and thus influence how any given subject responds, etc. (After all, the thought experiments being used are thought to be representative of a specific philosophical idea by philosophers themselves. It is plausible that they are interpreted in many different ways by laypeople who are not used to thinking about certain terms or phrases in the ways of Western analytic philosophy.) My hope with this dissertation was to investigate, question, and attempt to corroborate this ideal story by applying a layer of qualitative data over the top of what happens at the quantitative level during an experiment. The experiment I chose to investigate is OIC experiment 3, but others could be explored in this fashion too.

My strategy for investigation was to zoom in on this ideal story within OIC experiment 3 with the help of qualitative methodologies that are less restricted than the quantitative survey built into this experiment. First, I used the think aloud method to gather subjects' real-time verbalizations as they engaged with the thought experiments in OIC experiment 3. Second, I interviewed subjects in order to clarify any of their thoughts and experiences regarding the vignettes, their think aloud verbalizations, or any other questions and concerns they might have had. In total, what I found was limited corroboration on the ideal story as it pertained to the most philosophically interesting questions (those surrounding "ought"), a great deal of the story breaking down, and a number of interesting and nuanced philosophical judgments relating to OIC that participants couldn't have feasibly articulated solely in the quantitative questions they were presented with.

In the end, I suggested a number of ways that my qualitative data reveal more nuance in people's thinking regarding OIC; each suggestion points to different ideas that future research could explore in quantitative, qualitative, and analytic research. For example, I suggested that imaginative resistance, concreteness, moral remainders, and more seem to influence some people's thinking regarding "ought implies can" in unique and sometimes combined fashions. I ultimately found qualitative methods to be indispensable, which points to triangulation of methods as a fruitful path for experimental philosophers to pursue with regard to OIC and any other philosophical debate in ethics and otherwise. While my project was narrowly focused on the OIC study, it speaks to issues almost certainly relevant in other x-phi debates and it presents methodologies almost certainly beneficial to their exploration.

I hope the reader finds my work here to be a successful exploration of new approaches to experimental philosophy research. As an exploration, my project involved some amount of

calculated risk (that I think paid off): it could have turned out that I was misled into thinking that qualitative methodologies would be useful for x-phi, and I might have ended up concluding this project with little to say in favor of triangulation of methods in the x-phi context. Ultimately, though, I think I succeeded in showing both that triangulation of my particular methods of choice—the survey, the think aloud method, and the interview—is useful for x-phi and that triangulation of methods in general is useful to x-phi. By examining people’s philosophical judgments from different methodological perspectives, I believe I succeeded in showing that some judgments can be straightforwardly captured by quantitative survey questions while others are more nuanced and require examination from open-ended qualitative methods in addition to the responses typical of quantitative surveys. When combined, my methods painted a clearer picture of how a participant thinks about obligation, blame, and ability than any one method could paint on its own. At the very least, I hope the reader agrees that the data I produced from think aloud sessions and interviews cast significant doubt on the positive claims made in x-phi on behalf of surveys alone; this positions qualitative methodologies to be used in x-phi in at least the one role of being a means of checking the accuracy of conclusions made on behalf of quantitative surveys. Since this project was exploratory, I can’t offer a fully developed answer to the question of how triangulation of methods should best operate in the x-phi context. This answer can only come from more experiments and more analysis of the epistemological and methodological frameworks that allow for the effective integration of data from multiple methodologies as they intend to capture people’s philosophical judgments. What I think I do answer, though, is whether combining multiple methodologies, both quantitative and qualitative in nature, can better inform our experimental philosophy. The answer is yes. I hope to be involved in future research that further develops the triangulation of methods in x-phi. There is much to explore by applying the methods I used to

different philosophical issues in ethics and beyond. And there is so much to experiment with by using different methodologies altogether. From here, the expanded x-phi toolbox looks more limitless than limited!

Through my passion and hard work, I hope to have unearthed something valuable for current philosophical conversation in moral philosophy as well as future philosophical research of the experimental and non-experimental kinds.

APPENDICES A THROUGH N

Appendix A: Warm-Up Survey 1, “Lying”^{504*}

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix B: Warm-Up Survey 2, “Giving”*

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix C: Main Survey 1, “Moral Obligation, High Fault, Unable”^{505*}

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix D: Main Survey 2, “Moral Obligation, Low Fault, Unable”^{506*}

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix E: Follow-Up Survey^{507*}

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix F: Recruitment Flyer*

*Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix G: Body of Recruitment Email

PARTICIPATE IN A PAID RESEARCH STUDY

- **Pays \$25 (short surveys & one interview, 35 - 50 minutes total)**
 - **Must be 18 years or older and native English speaker**

Location: Claremont Graduate University

*To sign up or get more info, email me:
kyle.thompson@cgu.edu*

Appendix H: Quantitative Data from My Study*

* Please see additional appendices documents for this appendix. For those appendices that have a format different than the main dissertation document, I needed to include them separately.

Appendix I: Researcher Script Version 1

- 1. Prior to giving a participant the informed consent form, I will read the following aloud:** “Welcome, and thank you so much for participating in this experiment. Here is your initial payment of \$5 for coming out for the experiment; you will receive the remaining \$20 after completing the experiment. Before we begin, I will ask that you read and sign an informed consent form. The form details what you will be asked to do in the study along with other important information. Please read it carefully and let me know if you have any questions.”
- 2. Prior to giving a participant the first Warm-Up Vignette, I will read the following aloud:** “There are three main parts of this experiment. The first involves you completing a survey while verbalizing your thoughts as you complete it. The second involves an interview, and the third involves an additional survey that you will complete without verbalizing your thoughts. Before we begin this first part, it is helpful to practice ‘thinking aloud.’ Here is a brief Warm-Up survey for you to complete while thinking aloud. When I say ‘begin,’ please start reading aloud from the top of this page and please make sure to say out loud any thoughts you have while reading and completing the survey. In short, you

are being asked to ‘think aloud’ and talk your way through this survey. I will not be recording the audio from this Warm-Up exercise unless you request that I do so. OK, and ‘begin.’”

3. **If it seems helpful for a participant to complete an additional Warm-Up vignette, I will suggest doing so and I will read the following out loud:** “I think it would be helpful to complete one more Warm-Up survey. Again, when I say ‘begin,’ please start from the top of this page, making sure to read aloud and verbalize your thoughts while you complete the survey. As with before, I will not be recording the audio from this Warm-Up exercise unless you request that I do so. OK, and ‘begin.’”
4. **Prior to giving a participant the first survey that they will complete while I am recording their thinking aloud, I will read the following aloud:** “We will now begin the first survey and I will be recording the audio as you complete it. When I say ‘begin,’ please start from the top of the page, reading out loud and thinking out loud just as you did in the Warm-Up session. OK, and ‘begin.’”
5. **Prior to conducting the interview, I will read the following out loud:** “Next, we will begin the interview, which will focus on the experience you had with the survey you just finished. I will be recording the audio from this interview. Please feel free to look at your completed survey during your interview in case it helps jog your memory of your responses.”
6. **Prior to giving a participant the second survey (follow-up survey), I will read the following aloud:** “Here is the second survey for you to complete. I will not be recording any more audio for the experiment, and you can simply fill out this survey.”
7. **After the surveys and interview are complete, I will read the following aloud:** “Thank you so much for your time and participation. Please let me know if you have any questions. You can also email me at kyle.thompson@cgu.edu if you think of a question later on. Lastly, if you know anyone who might be interested in participation in my study, please provide them with my email.”

Appendix J: Researcher Script Version 2

1. **Prior to giving a participant the informed consent form, I will read the following aloud:** “Welcome, and thank you so much for participating in this experiment. Here is your initial payment of \$5 for coming out for the experiment; you will receive the remaining \$20 after completing the experiment. Before we begin, I will ask that you read and sign an informed consent form. The form details what you will be asked to do in the study along with other important information. Please read it carefully and let me know if you have any questions.” [Make sure to offer participant copy of informed consent and water]
2. **Prior to giving a participant the first Warm-Up Vignette, I will read the following aloud:** “As mentioned on the informed consent form, I am interested to learn how people think about common philosophical ideas. There are three main parts of this experiment. The first involves you completing a survey while verbalizing your thoughts as you complete it. The second involves an interview, and the third involves an additional survey that you will complete without verbalizing your thoughts. Before we begin this first part, it is helpful to practice ‘thinking aloud.’ The reason I am asking you to think aloud is that

I am simply interested in *how* you think about the scenarios presented on the surveys. To practice, I will give you a brief Warm-Up survey for you to complete while thinking aloud. When I say ‘begin,’ please start reading aloud from the top of this page and please make sure to say out loud any thoughts you have while reading and completing the survey. In short, you are being asked to ‘think aloud’ and talk your way through this survey. I will not be recording the audio from this Warm-Up exercise unless you request that I do so. OK, and ‘begin.’”

3. **If it seems helpful for a participant to complete an additional Warm-Up vignette, I will suggest doing so and I will read the following out loud:** “I have one more Warm-Up survey for you to complete. Again, when I say ‘begin,’ please start from the top of this page, making sure to read aloud and verbalize your thoughts while you complete the survey. As with before, I will not be recording the audio from this Warm-Up exercise unless you request that I do so. OK, and ‘begin.’”
4. **Prior to giving a participant the first survey that they will complete while I am recording their thinking aloud, I will read the following aloud:** “We will now begin the first survey and I will be recording the audio as you complete it. When I say ‘begin,’ please start from the top of the page, reading out loud and thinking out loud just as you did in the Warm-Up session. OK, and ‘begin.’”
5. **Prior to conducting the interview, I will read the following out loud:** “Next, we will begin the interview, which will focus on the experience you had with the survey you just finished. I will be recording the audio from this interview. In this interview, I would like to learn more about your thinking as you read the scenario and completed the survey. Feel free to look at the survey or read from the survey during this interview. In fact, some of my questions will ask you about the survey itself, so it will be helpful for you to look at the survey during the interview.”
6. **Prior to giving a participant the second survey (follow-up survey), I will read the following aloud:** “Here is the second survey for you to complete. I will not be recording any more audio for the experiment, and you can simply fill out this survey without talking aloud.”
7. **After the surveys and interview are complete, I will read the following aloud:** “Thank you so much for your time and participation. Here is the remaining \$20 of payment. Please let me know if you have any questions. You can also email me at kyle.thompson@cgu.edu if you think of a question later on. Lastly, if you know anyone who might be interested in participation in my study, please provide them with my email.”

Appendix K: Representative Qualitative Interview Questions Version 1

Note: Participants will have access to their own surveys when being interviewed so they can reference the questions and their responses.

While some of the interview questions will be improvised based on a participant’s response to a certain question, the following questions represent a general outline of the interview that I will conduct (this list is not exact or exhaustive; it represents the kinds of questions I will ask):

Question I: I would like to learn more about your thinking as you read the scenario and completed the survey. Feel free to look at the survey or read from the survey during this interview.

- a. To begin, please tell me about your *first* reaction to the scenario you read about with Adams and Brown.
- b. What did **question 1** on the survey ask?
- c. What did you choose for **question 1** and why did you choose that response?
- d. What did **question 2** on the survey ask?
- e. What did you choose for **question 2** and why did you choose that response?
- f. What did **question 3** on the survey ask?
- g. What did you choose for **question 3** and why did you choose that response?
- h. **Question 1** asks about Brown at a specific time during the scenario: 5:50. However **question 2** doesn't ask about Brown at a specific time. Were you thinking of a specific time during the scenario when you answered **question 2**? Please explain.
- i. **Question 1** asks about Brown at a specific time during the scenario: 5:50. However **question 3** doesn't ask about Brown at a specific time. Were you thinking of a specific time during the scenario when you answered **question 3**? Please explain.
- j. What do you think about the following statement? "At 5:55, Brown ought to make it to the theater by 6."

Question II: Can you describe a time when you were in a situation similar to the one described in the scenario you read about?

Question III: The word "ought" is an important word in the survey. Please describe how you understood the word "ought" when it was used in **question 3**.

Question IV: Do you think the word "ought" can have different meanings? Please explain.

Question V: Are there things a person *ought* to do that he *can't* do? Please explain.

Question VI: Please describe something that you think you *ought* to do but that you are *unable* to do.

Question VII: Do you have any questions about the scenario or the three questions on the survey that you would like to ask me?

Question VIII: To close, I want to ask you about Brown and Adams.

- a. What do you think about Brown's actions? What would you say to him if he were here right now?
- b. And what do you think about Adams' actions? What would you say to him if he were here right now?

Appendix L: Representative Qualitative Interview Questions Version 2

Note: Participants will have access to their own surveys when being interviewed so they can reference the questions and their responses.

While some of the interview questions will be improvised based on a participant's response to a certain question, the following questions represent a general outline of the interview that I will conduct (this list is not exact nor exhaustive; it represents the kinds of questions I will ask):

Question I:

- k. To begin, please tell me about your *first* reaction to the scenario you read about with Adams and Brown.
- l. What did **question 1** on the survey ask?
- m. What did you choose for **question 1** and why did you choose that response?
- n. What did **question 2** on the survey ask?
- o. What did you choose for **question 2** and why did you choose that response?
- p. What did **question 3** on the survey ask?
- q. What did you choose for **question 3** and why did you choose that response?
- r. Do you agree or disagree with the following statement? "At 5:55, Brown ought to make it to the theater by 6." [*If the response is different than the response the participant gave on the survey for **question 3**, explore why this is. If it is the same or similar, ask the participant to explain his or her thinking.*]
- s. Do you agree or disagree with the following statement? "At 5:59, Brown ought to make it to the theater by 6." [*If the response is different than the response the participant gave on the survey for **question 3**, explore why this is. If it is the same or similar, ask the participant to explain his or her thinking.*]

- t. Do you agree or disagree with the following statement? “At 5:40, Brown is to blame for not making it to the theater by 6.” [*If the response is different than the response the participant gave on the survey for **question 2**, explore why this is. If it is the same or similar, ask the participant to explain his or her thinking.*]

Question II: Can you describe a time when you were in a situation similar to the one described in the scenario you read about?

Question III: [*If the participant selected “agree” for both question 1 and question 3, ask them part a and c; if the participant selected “agree” for question 1 and “disagree” for question 3, ask them part b and c; if neither is the case, then only ask them part c*]:

- a. For question 1 you agreed that at 5:50 Brown can’t make it to the theater by 6. Then, for question 3, you agreed that Brown ought to make it to the theater by 6. Looking at these two responses together, do you agree or disagree that Brown ought to do something that he can’t do?
- b. For question 1 you agreed that at 5:50 Brown can’t make it to the theater by 6. Then, for question 3, you disagreed that Brown ought to make it to the theater by 6. Looking at these two responses together, do you agree or disagree that Brown ought to do something that he can’t do?
- c. Are there things a person *ought* to do that he *can’t* do? Please explain.

Question IV: The word “ought” is an important word in the survey. Please describe how you understood the word “ought” when it was used in **question 3**.

Question V: Do you think the word “ought” can have different meanings? Please explain.

Question VI: Please describe something that you think you *ought* to do but that you are *unable* to do.

Question VII: What do you think about Brown’s actions? What would you say to him if he were here right now?

Question VIII: The scenario does not mention cell phones, and it does not mention whether or not Brown tries to contact Adams. If you knew that both Brown and Adams had cell phones to communicate on, would you view the scenario differently?

Question IX: To close, do you have any questions about the scenario or the three questions on the survey that you would like me to clarify?

Appendix M: On Transcription for My Study

What is a Transcript?⁵⁰⁸

Before looking at how I constructed my transcript, I think it is important to acknowledge what a transcript *is* with relation to the events that are transcribed in it. In *InterViews: An Introduction to Qualitative Research Interviewing*, Steinar Kvale makes some important points about transcription and transcripts that I find useful in describing my own transcription process.

In a chapter on transcribing, Kvale claims that “rather than being a simple clerical task, transcription is itself an interpretive process.”⁵⁰⁹ This is certainly a claim I appreciated to a degree going *into* my transcription—my first one ever—and one I appreciate immensely now having transcribed 40 think aloud sessions and interviews. The naïve idea that transcribing is simply a matter of focused typing is fundamentally wrong, as transcription involves making thousands of interpretative judgments in converting an audio recording to a text that appears to be more or less fixed. Simply put, Kvale is right to make this point because it highlights that the distinctiveness of the original interview being recorded in comparison to the resultant transcript. To this point, Kvale rightly argues that “transcripts are not copies or representations of some original reality, they are interpretative constructions that are useful tools for given purposes. Transcripts are decontextualized conversations, they are abstractions, as topographical maps are abstractions from the original landscape from which they are derived.”⁵¹⁰ And just as it would be a mistake to treat a map as the landscape, so too would it be to treat the transcript as the interview. Of course, this doesn’t negate the precision of a map or transcript in helping us understand the landscape or interview, respectively.

Kvale continues, then, by dismissing the notion that there exists a true or objective version of a given transcript: “the question ‘What is the correct transcription?’ cannot be answered—there is no true, objective transformation from the oral to the written mode.”⁵¹¹ Instead he notes it better to ask “‘What is a useful transcription for my research purposes?’”⁵¹² That is, because to transcribe is to interpret, then the text form of an interview is best tailored to the needs of a research project rather to pursuit of some misguided notion that there is an objectively correct version of a transcript waiting to be summoned from an audio recording. Real-time, oral conversation simply isn’t written text, so in the process of transcribing one thing to the other, we must accept something is lost in translation, or transformation.⁵¹³ What we have in a transcript is flattened; it is useful, but it is flattened.

So, what I have below is an overview of my transcription methods, those suited for my research needs. I acknowledge that my transcripts of the think aloud sessions and interviews aren’t identical to the oral conversations they detail, but I did make a strong effort to capture as much nuance as possible without transcribing things that I could very easily misinterpret in the transcription process. Kvale argues that “there is one basic rule in transcription—state explicitly in the report how the transcriptions were made.”⁵¹⁴ In alignment with this, I have methodically detailed my methods below.

An Overview of How I Constructed My Transcripts

My overriding goal in transcribing was to produce verbatim transcripts that represent participants’ utterances without actively attempting to capture any affective feature of their utterances (e.g., whether they uttered something sympathetically), any features of their intonation

as they spoke (e.g., whether they modulated their pitch as they spoke), or other features of their tone of voice they used in speaking (e.g., whether they uttered something with sarcasm). However, while I didn't actively seek to transcribe affective features of utterances, intonation, or tone of voice, I no doubt captured such features at times, especially when I decided to use this punctuation over that one; I explore the specifics of my punctuation choices below. In sum, then, my main goal was to capture the words people were uttering without attempting to interpret complex features of the *way* in which those words were uttered.

I aimed for this verbatim *style* from the beginning, and I made certain judgment calls on more technical features of the transcript as I encountered new situations in the audio. For example, when I encountered laughter, I made a decision as to whether or not I would transcribe laughter whenever I encountered it again for the remainder of the transcribing process. I did not enter into the transcribing process with a pre-determined list of specific rules that I would employ. Rather, I entered in with certain general principles in mind—e.g., that I wanted to represent every speech act in the transcript—and then made a few important judgments once I encountered specific moments in the audio recordings. My judgments on these situations then became precedent for future instances, thereby becoming codified rules (detailed below).

I intentionally refer to transcription *style* rather than transcription methodology, though my style is informed by method, because I think the idiosyncrasies that I possess, or that any other transcriber possesses, should not benefit from the air of objectivity imbued by the term “methodology.” There are many small judgments that one makes as a transcriber—e.g., whether the utterance heard was “want to” or more accurately “wanna,” or whether this or that punctuation accurately reflects a speaker's vocal rhythms. Even if one defers to a methodological principle—e.g., *use a question mark only when it is clear that a participant is asking a question*—it still remains a judgment call as to what is or isn't an instance in which the principle applies. What this amounts to, then, is a *style* of transcribing not dissimilar to how we speak of a writer's style, though the former is more constrained than the latter since it is bound by an audio recording rather than versatile grammatical conventions alone. So, while I was principled in my transcribing, my transcribing style is somewhat determined by my idiosyncrasies and a few judgment calls I made for reasons I want to explain.

Here is a list of rules I employed after encountering a first instance; I give reasons for my decisions, too:

1. **I transcribed filler words.** Such expressions could display that a person is pondering something, or that she is unable to articulate an idea without stopping to correct it, or so many other things. An “um” that I write with three m's isn't necessary longer in actual speech than one that I transcribed as having one or two m's. The number of m's is largely an immediate, stylistic choice: what I initially typed for a given filler word upon first hearing it within the context of the think aloud session or the interview was what I ended up using. Also, I sometimes had certain utterances copied and ready to paste for quicker transcription. For example, in some cases I had “umm” copied so I could easily paste it into a document time and time again when I felt that “umm” would appropriately capture what was heard on the audio recordings. Even then, with copied and pasted material, I would always edit it if the pasted item didn't fit what I thought the transcript needed; it just helped me quickly convey “umm” after “umm” for those utterances that would be properly captured by an “um” with two m's.

2. **I transcribed interruptions.** Interruptions are very much a part of conversations and could potentially reveal important information as to why a given person responds the way he or she does.
3. **I transcribed pauses.** Pauses might show that a participant took time to reconsider something, which is why they are salient to my study. The pauses I report weren't timed, though I did my best to mark whether I thought a pause was brief, in which case I simply wrote "pause," or significantly longer, in which case I specified that it was a longer pause.
4. **I transcribed laughter.** I did so for similar reasons as filler words and interruptions.
5. **I transcribed whispering when it appeared especially blatant and noteworthy to me.** I did so for similar reasons as filler words and interruptions.
6. **I transcribed sighs.** I did so for similar reasons as filler words and interruptions.
7. **I didn't transcribe coughs or similar guttural noises.** I didn't transcribe these noises because it is sometimes difficult to know who is actually coughing, me or the interviewee. Also, coughing is largely irrelevant to the think aloud session or the interview given the mundane nature of the conversations I was having about moral obligation. In another context, such as a stressful interrogation, an interviewee's coughing might be thought to be caused by a nervousness that calls into question the interviewee's response, or coughing might be seen as a tell that betrays an interviewee's dishonesty. In my study, though, there is no reason to think that guttural noises are indicative of special cognitive processing; they almost certainly result from a need for physical relief. And, even if such noises for this or that participant were relevant to their thinking, there would be no reliable way to discern this from the lifeless text of the transcript. Therefore, I think my decision to abstain from transcribing coughing and the like best serves the kind of philosophical analysis I performed in this project.
8. **I didn't transcribe pen clicking, audible writing on paper, or water drinking.** Who is clicking the pen? What is the participant scribbling in that particular moment? Who is drinking the water? I can't answer these questions across all examples, so I opted not to transcribe them. I also feared that transcribing some of these noises might end up making them seem more prominent in the interview experience than they were: a microphone picks up and picks out very subtle noises that when listened to in headphones sound much more vivid and salient to a soundscape than when they actually occurred. If one reads in a transcript that a participant is scribbling and clicking her pen here and there, one might think these are especially important noises in the global interview experience *because* I went to the pains of transcribing them. In actuality, they are likely background noises that went unnoticed at the time of the interview. Therefore, for fear of making much out of sounds I can't confidently interpret, I excluded such noises in the transcript.
9. **I didn't transcribe the affective features of participants' utterances, features of their intonation, or features of their tone of voice (except in cases in which it indirectly influenced my choice of punctuation, such as inputting question marks).** I don't feel I have the expertise to accurately interpret and record such features of a person's speech, so I opted not to do so.

If the above list details what I did and didn't attend to, the next list details how I textually denoted specific features of my transcript:

1. **Dots.** I use dots (...) *not* to indicate long pauses, but to indicate small breaks in sentences or to show continuation between verbalizations that were relatively connected and fluid. In other words, the ellipses (or dots) function as connectors between phrases, terms, utterances, re-statements, and more.
2. **Brackets.** Words in brackets—e.g., [laughs]—apply to the person talking unless otherwise noted. If words in brackets are not in quotation marks, then they detail what happened: i.e., that the person laughs. If the transcript says the word “unintelligible”⁵¹⁵ in brackets—i.e., [unintelligible]—then the person is saying something that I couldn’t understand from the audio recording for whatever reason. If words in brackets are in quotation marks, then they refer to what the person says, and in these cases I will specify who is saying something if it is not the current speaker.
3. **Interruptions.** I always made sure to represent the words and phrases uttered by both parties, even when an interruption was made, but I transcribed certain interruptions using brackets—i.e., I would insert the interrupters words in brackets and indicate an interruption in the midst of the text of the interrupted party—rather than starting a new line of dialogue. In other cases, I opted for a new line of dialogue. In each case, the transcript clearly articulates who is interrupting whom and what is being said by each party, but I didn’t decide on just one rule for transcribing these instances because I treated each situation as somewhat unique and transcribed the interruption as I thought best articulated a given moment in the audio.
4. **Placement of Laughter and Interruptions.** Laughter and interruptions were transcribed similarly: I tried to put the brackets exactly where the interruption or laughter took place, but the brackets more so indicate the *near-precise vicinity* of the laughter (certain words might fall just before or after the bracketed notation because it is very difficult to hear exactly when laughter begins or an interruption occurs). I simply don’t want to claim to be so exact as to make claims based on the atomically precise moment of interruption or laughter. That said, I want to express that I often listened to portions of audio multiple times in order verify that the transcript I produced was accurate.
5. **General Punctuation.** I avoided using too much punctuation⁵¹⁶, but I often did include periods and question marks where they seemed appropriate. Many sentences that I marked with a period probably were questions but lacked a clear intonation of one so I opted for a period. These can be judged by the reader based on the surrounding textual context.
6. **Emphasis.** I italicized words that were particularly emphasized in the audio; this rarely occurred.
7. **Unintelligible audio.** Something marked “unintelligible” denotes that the very spot marked as such is the unintelligible bit.

8. **Sanitized information.** I replaced potentially identifying information with invented information. I erred on the side of being extra careful in cases where information could be identifying.
9. **Template-based text.** Because I asked participants to read the surveys aloud, and because I generally followed an interview script, I copied and pasted a template into some (if not all) of the transcripts for transcribing convenience. Of course, I edited them extensively based on the audio I listened to since I didn't always read my questions exactly the same and participants often omitted words or added them and so on. I just want to note this, though, for the sake of transparency.

Lastly, I want to acknowledge that readers of my transcripts will likely notice many moments in the transcripts that appear as if they could be errors in my transcription just as easily as they could be atypical utterances from participants. In some cases, it might even seem obvious that I, the transcriber, has erred, and that surely the speaker didn't say what is written, because it is atypically worded or somewhat incoherent. While I can't guarantee that I never mistyped during my transcribing, I am confident that my transcribing is accurate and that what one reads really is what was spoken; I am confident because I listened to moments in the audio recordings multiple times to ensure accuracy. Had I opted to label particular moments using "*sic*," the transcript would have been flooded with the label; I was surprised to learn how often people's utterances contain small or even large "errors" which might go unnoticed when they are being heard in real time. As detailed in the informed consent form participants signed, I deleted the audio files once they were transcribed, so I can't return to check an additional time any atypical utterance in the transcript. While it is possible that I erred during transcription and checking my work, I am confident in treating the transcript as accurate in those moments when a speaker said something that reads as a bit awkward or out of place. In effect, the reader is advised to put a blanket "*sic*" over all my transcripts.

Appendix N:

Additional Information on My Data

Including Trivial Coding and Analysis Choices

In an effort to be as transparent as possible in how I report my data, I've included this appendix. This information is largely trivial since the few coding choices and analysis choices I made are based in common sense. This information, in short, helps the reader understand how I arrived at the data in Appendix H and my statistical analyses in chapter five for those minute details that aren't worth full discussion in the body of my dissertation.

1. First, I want to clarify that I didn't record *which order* I administered the two warm-up surveys, "Lying" and "Giving," whenever both were administered, which was to everyone but participant 1. If memory serves me right, I tended to give participants "Lying" then "Giving" as the former is shorter. However, I didn't record the order for any given participant. What I did record was whether they did just one warm-up or two warm-ups and which one they completed if it was just one. In just one case, participant 1, I gave just "Lying."
2. Next, I want to explain the few "switched" scores detailed on the second tab of Appendix H as opposed to the "original" scores detailed on the first tab. For questions 1-3 on the main survey, the *original score* is always what the participant had marked on the Main Survey 1 or Main Survey 2 during the think aloud session as their final and definitive response during the survey completion. That is, the original response, the one reported on the first tab of Appendix H would be each subject's response were he or she to have *not* completed the interview. As for those scores that are switched, the participant switched his or her response after the think aloud session and *during* the interview session.
3. In the end, I opted not to use the *switched* scores for data analysis for one main reason I'd like to discuss. It is tempting to statistically analyze my switched data, especially in light of the fact that at least one participant, participant 12, switched her response on question 1 upon realizing she read the scale incorrectly. In a sense, I could just "see" what the data would look like if she responded to question 1 in alignment with her judgment regarding Brown's ability to make his meeting. But once I open up this kind of "corrective," I am assuming a great many things about the rest of my data as well as

Chituc et al.'s data to which I compare my own. That is, once I decide to make "corrections" then I need to assume that other data points *aren't* in need of correction, and that Chituc et al.'s data *don't* result from scale misreading. But even a brief glance at the Chituc et al. data makes it seem likely that some people misread the scale for their question about "can" (and probably other questions), too. In experiment 2 of the OIC study, participant 126, who agreed with a 50 that Adams *can* make his meeting, explains that "THE MANS CAR IS BROKEN DOWN THERE IS NO WAY THAT ADAMS IS CAPABLE OF MEETING WITH BROWN."⁵¹⁷ And again, participant 176 in experiment 2, who also agreed with a 50 that Adams *can* make the meeting, writes the following explanation: "He broke down, he cannot meet."⁵¹⁸ There might be another explanation here, but it is probable that these participants intended to express that Adams *can't* make it, given their qualitative explanations. And so, presumably some small percentage of people in OIC experiment 3 misread questions, too. Since not all of these kinds of "errors" can be fixed and attended to with certainty—after all, some people might have misread the scale and then failed to produce any qualitative data that would suggest that this is the case—then I think it best to let the quantitative data stand as is and instead highlight such discrepancies as they are relevant in the qualitative analysis. Ideally, any random errors will fall on both sides of a given quantitative scale and be washed out in the larger data trends.

4. Next, I want to clarify some small coding choices. As apparent on Appendix H, I coded the first question on the follow-up survey (Appendix E) with three different values relating to whether or not a participant had taken a philosophy course in the past. While this is a straightforward bit of coding, there is one judgment I made. One participant,

- participant 10, indicated some additional information regarding the completion of a philosophy course: “It was technically Religious Studies, but dealt w/ Religion on a philosophical level.” Participant 10 *wrote* this just as I have preserved it *in addition to* indicating a “Yes” on the survey. I chose to code this simply as a “Yes” that they have taken a philosophy course as this serves as the kind of philosophical training I am trying to note.
5. Again, as seen on Appendix H, I coded responses from the follow-up survey (Appendix E). The question about educational attainment, number 4, was straightforward enough, and involved just one bit of judgment. Participant 27 indicated he is attending college as a rising junior by writing “In College, Rising Junior.” I coded this as “some college” since the subject reported that he is in college as a rising junior which means he hasn't completed college.
 6. Here is another small coding judgment seen on Appendix H. The question about gender on the follow-up survey, 6, was a fill in the blank question, so I copied down what participants wrote. All of the responses didn't require a judgment on my part except participant 29, who wrote “Woman (Cis-Het)” (though the capital letters might not be from the original). I coded this participant as female.
 7. In addition, I want to note that in reviewing participant 2's transcript, I noticed that he talks about his selections on questions 1 and 2 as being “5” for each, when I recorded that he selected “2” for each. In an endnote in chapter six (number 329), I elaborate on this. In short, though, I think it most likely he misspoke in reference to his selections and that my recorded data are correct, but this of course might not be the case.

8. Lastly, I want to note that I didn't code the fill in the blank responses from the follow-up survey question about race/ethnicity, number 7. So, I copied down what participants wrote, though the capital letters might not be from the original in each case. So, for example, participant 4 put actually wrote the “?” seen in her response as reported on Appendix H.

¹ Joshua Knobe and Shaun Nichols, “An Experimental Philosophy Manifesto,” in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 8-9.

² The idea to use the think aloud method within an experimental philosophy study comes from philosopher James Andow. For specific details, see James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1130, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.

³ Kwame Anthony Appiah discusses x-phi and references how it views the armchair, depicted visually in a YouTube video entitled “Experimental Philosophy Anthem” where an armchair burns. See Kwame Anthony Appiah, “The New New Philosophy,” Idea Lab, *The New York Times Magazine*, December 9, 2007, <https://www.nytimes.com/2007/12/09/magazine/09wwln-idealab-t.html>, and also Alina Simone, *Experimental Philosophy Anthem*, performed by Alina Simone (2007), YouTube, <https://www.youtube.com/watch?v=tt5Kxv8eCTA>.

⁴ Immanuel Kant, *Critique of Pure Reason*, trans. Norman Kemp Smith (New York: Palgrave Macmillan, 2007), 473.

⁵ This definition of laypeople largely emerged from my reading of various definitions in the x-phi literature. I do draw some important features from a specific definition, one Justin Sytsma provides in the midst of a chapter on experimental philosophy of mind, that clarifies “laypeople” as “the ‘folk’ or people with little to no training in philosophy or consciousness studies.” In particular, I like that this definition makes special reference to “consciousness studies” which is important to the chapter; this gave me the idea to make special reference to “ought implies can” since that is the special focus of my work. For the quotation and further information, see Justin Sytsma, “The Robots of the Dawn of Experimental Philosophy of Mind,” in *Current Controversies in Experimental Philosophy*, edited by Elizabeth O’Neill and Eduard Machery (New York: Routledge, 2014), 50.

⁶ While I have opted to use “judgments” over “intuitions” as much as possible throughout this project, there might be traces of “intuition” language since I initially went into this dissertation thinking I was talking specifically and only about “intuitions.” It was only later that I realized, through conversation about my project, that I was concerned with “judgments” more generally.

⁷ Joshua Alexander. *Experimental Philosophy: An Introduction* (Cambridge: Polity Press, 2012), 11-27.

⁸ I arrived at this definition after reading some literature in management studies and experimental philosophy centered on “intuition” or “intuitions.” In particular, though, I derive the notion of intuitions being spontaneous from a particular definition by Weinberg, Nichols, and Stich (here they are referring to epistemic intuitions): “As we use the notion, an epistemic intuition is simply a spontaneous judgment about the epistemic properties of some specific case—a judgment for which the person making the judgment may be able to offer no plausible justification.” From Jonathan M. Weinberg, Shaun Nichols, and Stephen P. Stich, “Normativity and Epistemic Intuitions,” in *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 19. In addition, I derive the sense of an intuition being a “seeming” from Joshua Alexander’s examination of intuitions. In reference to a Gettier case, Alexander describes the *seeming* nature of intuitions: “Smith has the justified true belief that the man who will get the job has ten coins in his pocket. Yet, it doesn’t *seem* that he knows this.” From Joshua Alexander. *Experimental Philosophy: An Introduction* (Cambridge: Polity Press, 2012), 13.

⁹ Rhidian Hughes, “Vignettes,” in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 2. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n482>.

¹⁰ Immanuel Kant, *Critique of Pure Reason*, trans. Norman Kemp Smith (New York: Palgrave Macmillan, 2007), 473.

¹¹ Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 249-250. Also, for an expression of this contraposition, see Miklos Kurthy and Holly Lawford-Smith, “A Brief Note on the Ambiguity of ‘Ought’. Reply to Moti Mizrahi’s ‘Ought, Can, and Presupposition: An Experimental Study,’” *Methodes* 4, no. 6 (2015): 247.

¹² Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 250.

¹³ In conversation with Kurthy and Lawford-Smith, Mizrahi treats this inertness as saying that the claim of ought becomes “irrational” or “pointless” or “unfair” if an agent cannot do that thing. See Mizrahi, Moti, “Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith,” *Methodes* 4, no. 6 (2015): 252.

¹⁴ Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 250.

¹⁵ Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 256.

-
- ¹⁶ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2, (1984): 258.
- ¹⁷ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁸ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁹ Paul Henne et al., "An Empirical Refutation of 'Ought' Implies 'Can,'" *Analysis* 76, no. 3 (2016): 284.
- ²⁰ Paul Henne et al., "An Empirical Refutation of 'Ought' Implies 'Can,'" *Analysis* 76, no. 3 (2016): 284.
- ²¹ I like both of these descriptions—i.e., consistency and compatibility with OIC—and they are both employed and inspired by the work of Kurthy et al. See Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 8, no. 4 (2017): 1-24, <https://doi.org/10.1371/journal.pone.0175206>.
- ²² Mizrahi, Moti, "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith," *Method* 4, no. 6 (2015): 237.
- ²³ "X-phi" is used as an abbreviation for "experimental philosophy." See "about experimental philosophy," *experimental philosophy*, accessed June 2, 2017, <http://experimental-philosophy.yale.edu/ExperimentalPhilosophy.html>.
- ²⁴ I am alluding to the well-known fake barn case that Alvin I. Goldman discusses. For more details see Alvin I. Goldman, "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73, no. 20 (1976): 772-773.
- ²⁵ I am alluding to the famous trolley problem. For more details see Philippa Foot, "The Problem of Abortion and the Doctrine of the Double Effect," *Oxford Review* 5 (1967): 5-15, accessed June 5, 2017.
- ²⁶ "X-phi" is used as an abbreviation for "experimental philosophy." See "about experimental philosophy," *experimental philosophy*, accessed June 2, 2017, <http://experimental-philosophy.yale.edu/ExperimentalPhilosophy.html>.
- ²⁷ This definition results from my reading of and about experimental philosophy studies. Therefore, it is not a definitive definition, so to speak. Nonetheless, it coincides nicely, at least in part, with this definition from Joshua Alexander's introductory text: "In recent years, *experimental philosophy* has emerged as an exciting new approach to the study of people's philosophical intuitions. Experimental philosophers apply the methods of the social and cognitive sciences to the study of philosophical cognition since these methods are better suited than introspective methods to the study of what people, especially *other* people, actually think." From Joshua Alexander, *Experimental Philosophy: An Introduction* (Cambridge: Polity Press, 2012), 2.
- ²⁸ See term definitions in chapter one for more information on "judgments" vs. "intuitions."
- ²⁹ See term definitions in chapter one for more information.
- ³⁰ Joshua Knobe and Shaun Nichols, in "An Experimental Philosophy Manifesto," discuss how intuition and philosophical theory are sometimes related and how experimental philosophy can bear on this relationship. While they exclusively use the term "intuition" here, the same sentiment can be applied to what I am referring to as "judgment." In reference to experimental philosophy research, Knobe and Nichols write, "First we use the experimental results to develop a theory about the underlying psychological processes that generate people's intuitions; then we use our theory about the psychological processes to determine whether or not those intuitions are warranted." They continue, "Of course, this sort of question becomes especially pressing in cases where the intuitions are actually serving as evidence for a particular philosophical view." Joshua Knobe and Shaun Nichols, "An Experimental Philosophy Manifesto," in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 8.
- ³¹ Joshua Knobe and Shaun Nichols, "An Experimental Philosophy Manifesto," in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 7-11. Again, the authors here speak of "intuitions," but the same sentiment can be applied to what I am calling "judgments."
- ³² See Joshua Knobe, "Free Will and the Scientific Vision," in *Current Controversies in Experimental Philosophy*, edited by Elizabeth O'Neill and Eduard Machery (New York: Routledge, 2014), 69-85.
- ³³ I suppose some kind of citation might be necessary for my reference to "justified true belief" or "JTB" despite its prominence in philosophical literature. For a brief overview of the philosophical claims underlying JTB (and their supposed flaws), see Edmund L. Gettier "Is Justified True Belief Knowledge?" *Analysis* 23, no. 6 (1963): 121-123.
- ³⁴ Eddy Nahmias et al., "Is Incompatibilism Intuitive?" in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 81.
- ³⁵ Eddy Nahmias et al., "Is Incompatibilism Intuitive?" in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 81.

-
- ³⁶ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87.
- ³⁷ In the chapter I’ve been discussing, the experiment is presented just as a quantitative survey, but in the publication that describes the experiment in more detail (which the chapter points to in an endnote), there is mention of a qualitative component: subjects were given the opportunity to explain their responses on the backs of their surveys. And, in this more detailed publication, the authors note that they make changes to their pilot experimental design after qualitative explanations indicated that participants weren’t understanding the survey as intended. See Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 566, doi: 10.1080/09515080500264180.
- ³⁸ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87.
- ³⁹ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87.
- ⁴⁰ Yes, the data that I reference here is hypothetical based on my invented thought experiment about Paul. I feel confident asserting that most people would deem this to be an immoral act.
- ⁴¹ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 566, doi: 10.1080/09515080500264180.
- ⁴² Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 566, doi: 10.1080/09515080500264180.
- ⁴³ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 97.
- ⁴⁴ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 566, doi: 10.1080/09515080500264180.
- ⁴⁵ Shaun Nichols and Joshua Knobe, “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions,” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 109.
- ⁴⁶ Shaun Nichols and Joshua Knobe, “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions,” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 109.
- ⁴⁷ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 580, doi: 10.1080/09515080500264180.
- ⁴⁸ Tamar Szabó Gendler and Shen-yi Liao, “The Problem of Imaginative Resistance,” in *The Routledge Companion to Philosophy of Literature*, eds. Noël Carroll and John Gibson (New York: Routledge, 2015): 405-418, doi-org.ccl.idm.oclc.org/10.4324/9781315708935.
- ⁴⁹ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87-89.
- ⁵⁰ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 573-574, doi: 10.1080/09515080500264180.
- ⁵¹ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 97.
- ⁵² Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 580, doi: 10.1080/09515080500264180.
- ⁵³ Hagop Sarkissian et al., “Is Belief in Free Will a Cultural Universal?” *Mind & Language* 25, no. 3 (2010): 346-358, doi: 10.1111/j.1468-0017.2010.01393.x.
- ⁵⁴ Hagop Sarkissian et al., “Is Belief in Free Will a Cultural Universal?” *Mind & Language* 25, no. 3 (2010): 348, doi: 10.1111/j.1468-0017.2010.01393.x.
- ⁵⁵ Tamler Sommers, “Experimental Philosophy and Free Will,” *Philosophy Compass* 5, no. 2 (2010): 209, doi: 10.1111/j.1747-9991.2009.00273.x.
- ⁵⁶ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁵⁷ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁵⁸ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁵⁹ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.

-
- ⁶⁰ Robert Kirk, “Zombies,” *The Stanford Encyclopedia of Philosophy*, Spring 2019, <https://plato.stanford.edu/entries/zombies/>.
- ⁶¹ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁶² Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁶³ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87.
- ⁶⁴ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁶⁵ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁶⁶ I realize this reads a bit like question-begging: Because it hasn’t been tested by multiple methods then it isn’t valid, but what I mean by invalid is to be not tested by multiple methods. To avoid this circularity, though, I think there are both intuitive reasons—those that I’m slowly unpacking by providing other plausible explanations—and evidentiary reasons—those that will be forthcoming later in this project when I cite qualitative data that points to a lack of epistemic certainty with respect to quantitative data alone.
- ⁶⁷ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88-89.
- ⁶⁸ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 89.
- ⁶⁹ It is difficult to know if Feltz and Millan want their definition of “free will no matter what” to be tied so specifically to fatalism. The quote that follows is an instance in which they do tie the intuitions to fatalism, but the name “no matter what” and a moment on a page 551 wherein they connect such intuitions with cases of manipulation suggest they refer to literally what they are labelled as: “free will no matter what.” See Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 551. doi: 10.1080/09515089.2013.865513.
- ⁷⁰ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 549. doi: 10.1080/09515089.2013.865513. They earlier introduce these intuitions on page 533.
- ⁷¹ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 551. doi: 10.1080/09515089.2013.865513.
- ⁷² Of course, if Feltz and Millan don’t really want these intuitions to be defined in the context of fatalism as I suggest a few endnotes earlier (in endnote number 69), then my thinking is very similar (but not identical) to their thinking.
- ⁷³ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 530. doi: 10.1080/09515089.2013.865513.
- ⁷⁴ James Andow and Florian Cova, “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan,” *Philosophical Psychology* 29, no. 4 (2016): 550-566. doi: 10.1080/09515089.2015.1082542.
- ⁷⁵ James Andow and Florian Cova, “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan,” *Philosophical Psychology* 29, no. 4 (2016): 564. doi: 10.1080/09515089.2015.1082542.
- ⁷⁶ James Andow and Florian Cova, “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan,” *Philosophical Psychology* 29, no. 4 (2016): 564. doi: 10.1080/09515089.2015.1082542.
- ⁷⁷ James Andow and Florian Cova, “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan,” *Philosophical Psychology* 29, no. 4 (2016): 562. doi: 10.1080/09515089.2015.1082542.
- ⁷⁸ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 546. doi: 10.1080/09515089.2013.865513.
- ⁷⁹ James Andow and Florian Cova, “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan,” *Philosophical Psychology* 29, no. 4 (2016): 562. doi: 10.1080/09515089.2015.1082542.
- ⁸⁰ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 551. doi: 10.1080/09515089.2013.865513. Here, Feltz and Millan reference the work of Cokely and Kelley regarding the value of protocol analysis; see Edward T. Cokely and Colleen M. Kelley, “Cognitive Abilities and Superior Decision Making Under Risk: A Protocol Analysis and Process Model Evaluation,” *Judgment and Decision Making* 4, no. 1 (2009): 20-33.
- ⁸¹ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 533. doi: 10.1080/09515089.2013.865513.

-
- ⁸² Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 91.
- ⁸³ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 572, doi: 10.1080/09515080500264180.
- ⁸⁴ Adam Feltz and Melissa Millan, “An Error Theory for Compatibilist Intuitions,” *Philosophical Psychology* 28, no. 4 (2015): 551. doi: 10.1080/09515089.2013.865513.
- ⁸⁵ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 89.
- ⁸⁶ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 97.
- ⁸⁷ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 97.
- ⁸⁸ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 571, doi: 10.1080/09515080500264180.
- ⁸⁹ Dale Tuggy, “Trinity,” *The Stanford Encyclopedia of Philosophy*, Winter 2016, <https://plato.stanford.edu/entries/trinity/>. The doctrine of the Trinity, in its history and underlying meaning, is certainly more complex than this statement conveys, and I do not wish to make it seem so straightforward. In fact, in this entry, Tuggy explores some of the complexity of the analytic philosophy of the doctrine. For my purposes, though, this brief definition serves since I am just interested in grounding the constructed vignette that follows.
- ⁹⁰ Thanks to Eric Schwitzgebel for helping me develop this example; in particular his idea of a supercomputer spitting out an incoherent statement of a kind started the conversation leading to the Trinity case.
- ⁹¹ Tamar Szabó Gendler and Shen-yi Liao, “The Problem of Imaginative Resistance,” in *The Routledge Companion to Philosophy of Literature*, eds. Noël Carroll and John Gibson (New York: Routledge, 2015): 405, doi-org.ccl.idm.oclc.org/10.4324/9781315708935. In some straightforward sense, this quoted definition connects with what I am talking about in the Trinity case and people’s potential reactions to it. And, for now, I’m really only concerned with this general definition. However, it is worth noting that a good portion of this piece by Gendler and Liao, as well as other works on imaginative resistance, focuses on imaginative resistance in connection to morally objectionable or morally impossible material. I find myself confused when reading about imaginative resistance because it so often seems like the general definition is quickly shuffled away for some more focused look at scenarios wherein a person is asked to imagine something morally reprehensible, objectionable, etc. I discuss this more in chapter seven of this dissertation. If the canonical example of imaginative resistance is one of morally objectionable fiction, then I wonder if I am finding connections to imaginative resistance as a phenomenon that don’t exist when I invoke imaginative resistance to speak about resisting the confines of a supercomputer, for example.
- ⁹² Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 87.
- ⁹³ Eddy Nahmias et al., “Is Incompatibilism Intuitive?” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 88.
- ⁹⁴ Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology* 18, no. 5 (2005): 566-572. doi: 10.1080/09515080500264180.
- ⁹⁵ Tamler Sommers, “Experimental Philosophy and Free Will,” *Philosophy Compass* 5, no. 2 (2010): 204, doi: 10.1111/j.1747-9991.2009.00273.x.
- ⁹⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ⁹⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ⁹⁸ For an overview of similarly structured experiments, see Joshua Knobe and Shaun Nichols, eds., *Experimental Philosophy* (New York: Oxford University Press, 2008).
- ⁹⁹ Vlad Chituc and Paul Henne, “The Data Against Kant,” *New York Times*, February 19, 2016.
- ¹⁰⁰ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁰¹ This attribution to Kant explains the title of the New York Times article “The Data Against Kant.”

-
- ¹⁰² Immanuel Kant, *Critique of Pure Reason*, trans. Norman Kemp Smith (New York: Palgrave Macmillan, 2007), 473.
- ¹⁰³ For more information on precisely what relationship between “ought” and “can” is being conveyed when I refer to OIC in general and in conjunction with this study, see my extended terminological definition of OIC in chapter one. In short, though, entailment is what is at stake in their references and mine (unless otherwise specified).
- ¹⁰⁴ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁰⁵ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁰⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁰⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹⁰⁸ Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ¹⁰⁹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21-22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁰ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹¹ I have yet to come across such disclaimers in other studies either.
- ¹¹² This new and improved ideal story comes from Eric Schwitzgebel who provided feedback on my original. This new structure and its core components are of his design. Also, I want to note that the vignettes I’m referencing as examples in this ideal story are essentially those from OIC experiment 3.
- ¹¹³ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁴ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁵ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. I won’t make much of it here, but this language of “still” is potentially problematic in the way that it might bias an interpretation of the scenario such that participants feel, in the word “still,” that they are being asked to view obligation as something that could go away—to ask if you are still hungry implies that hunger is the kind of thing that can and does go away. Fortunately, OIC experiment 3, which I focus most on, abandons this language.
- ¹¹⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁸ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ¹¹⁹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

¹⁵⁰ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵¹ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵² Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

¹⁵³ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

¹⁵⁴ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵⁵ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵⁶ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵⁷ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵⁸ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁵⁹ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁶⁰ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁶¹ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See participant in the data with failed checks removed tab.

¹⁶² Again, I am making reference to the fake barn case that Alvin I. Goldman discusses; the very specifics of his case are different but the general point and thought experiment structure is held intact in my description here. For more details see Alvin I. Goldman, “Discrimination and Perceptual Knowledge,” *The Journal of Philosophy* 73, no. 20 (1976): 772-773.

¹⁶³ Edmund L. Gettier, “Is Justified True Belief Knowledge?” *Analysis* 23, no. 6 (1963): 121-123.

¹⁶⁴ To give a sense of the far-reaching nature of Gettier cases and the so-called Gettier problem, I am referencing the PhilPapers section dedicated to it in which Ian M. Church briefly describes how Gettier’s work has spawned a substantial debate within philosophy. It is also worth noting that PhilPapers has a section of its library dedicated to the Gettier Problem. See “The Gettier Problem,” PhilPapers, ed. Ian M. Church, accessed May 11, 2018, <https://philpapers.org/browse/the-gettier-problem>.

- ¹⁶⁵ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 1-4. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁶⁶ I've chosen to focus on triangulation over a similar and somewhat overlapping practice of mixed methods research (see Creswell source at end of this note) because the former, to my understanding, can be more expansive than the latter. My specific project certainly qualifies as a mixed methods study given that it mixes qualitative methods with quantitative ones, a key feature of mixed methods, but future x-phi research might seek to triangulate on people's judgments by making sole use of qualitative or quantitative approaches, or by use of methodologies that don't clearly fit in with the mixed methods paradigm. My emphasis on triangulation, and my borrowing of the concept from social science research, is to say that x-phi would benefit from a much more expansive and creative use of research methodologies rather than a narrow appeal to quantitative surveys as revealing of philosophical judgments. The central idea in triangulation (see Rothbauer source mentioned at end of this note) is one of coming to better understand a phenomenon by approaching it from different methodological perspectives; that is what interests me, not some call for making sure to always mix qualitative with quantitative methods. See John W. Creswell, "Mixed Methods Research," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 1-5. <http://dx.doi.org/10.4135/9781412963909.n269>. For more on triangulation, see Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 1-4. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁶⁷ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 2. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁶⁸ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 2. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁶⁹ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 2-4. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁷⁰ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 3-4.
- ¹⁷¹ Paulette M. Rothbauer, "Triangulation," in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 2. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- ¹⁷² Louise K. Thornton et al. "Reasons for Substance Use Among People With Psychotic Disorders: Method Triangulation Approach," *Psychology of Addictive Behaviors* 26, no. 2 (2012): 279-288, doi: 10.1037/a0026469.
- ¹⁷³ Louise K. Thornton et al. "Reasons for Substance Use Among People With Psychotic Disorders: Method Triangulation Approach," *Psychology of Addictive Behaviors* 26, no. 2 (2012): 287, doi: 10.1037/a0026469.
- ¹⁷⁴ Ronald Inglehart, Miguel Basañez, and Alejandro Moreno, *Human Values and Beliefs: A Cross-Cultural Sourcebook* (Ann Arbor: University of Michigan Press, 1998), 251. ebook.
- ¹⁷⁵ Ronald Inglehart, Miguel Basañez, and Alejandro Moreno, *Human Values and Beliefs: A Cross-Cultural Sourcebook* (Ann Arbor: University of Michigan Press, 1998), 251. ebook.
- ¹⁷⁶ "Denmark: Percent Christian," *The Global Economy*, accessed April 24, 2019, <https://www.theglobaleconomy.com/Denmark/christians/>.
- ¹⁷⁷ "The World Factbook: Denmark," *Central Intelligence Agency: Library*, last updated April 3, 2019, <https://www.cia.gov/library/publications/the-world-factbook/geos/da.html>.
- ¹⁷⁸ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 10-11.
- ¹⁷⁹ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 11.
- ¹⁸⁰ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 11.
- ¹⁸¹ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 11.
- ¹⁸² Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 11.
- ¹⁸³ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About*

Contentment (New York: New York University Press, 2008), 157.

¹⁸⁴ Phil Zuckerman, *Society Without God: What the Least Religious Nations Can Tell Us About Contentment* (New York: New York University Press, 2008), 157.

¹⁸⁵ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 26-30, doi: 10.1136/medethics-2015-102738.

¹⁸⁶ Because I will refer to this comment on multiple occasions, I want to label it as the *x-phi database comment*.

There are three important caveats I want to make in this moment and all future moments in which I make some kind of claim about how there are only *x* many studies in *x-phi* that use this or that method or something of that nature. First, many months have passed since I verified such claims, so it is probable that a few new studies have been published that negate my claim in some sense. This isn’t a concern for me since the point I want to make is mostly that there isn’t a large body of literature using this or that method. Second, it is entirely possible that there are studies that bear on *x-phi* investigations or would even qualify as *x-phi* investigations that are left out of my research. For example, Nahmias et al.—in their publication “The Phenomenology of Free Will”—reference work by psychologists published decades ago that could be categorized as *x-phi* in one way or another even if they don’t show up in this or that search of the relevant “*x-phi*” literature. It is beyond the scope of my project to look into works like these to see if they show up in a given database search of *x-phi* work and if they might be properly considered as work in *x-phi* if they don’t. Third, I have discovered that certain database searches just aren’t that great when compared to something like Google. I came across the Nahmias et al. study that I just referenced by following endnotes in another work by them and there I discovered that it actually contains use, albeit minimal use, of the talk-aloud method (which is described as identical or nearly so with the think aloud method) and a kind of interview process. Importantly, this study seems to count as work in *x-phi* and even shows up under in the Experimental Philosophy database of PhilPapers; however I couldn’t get it to appear in the database when I searched for terms like “talk aloud” and “talk-aloud.” This means that there might be more studies that slipped through my searches; it is still clear though that there isn’t a great deal of them if there are more, though. That all said, as I describe later, this Nahmias et al. study’s use of these methods is somewhat incomplete and discussed sparingly in the publication such that it doesn’t really act as an *x-phi* study that makes full use of talk-aloud or interview methods. This again points back to my general sentiment that there isn’t much in the way of these methods in full use as I used them in my study. See Eddy Nahmias et al., “The Phenomenology of Free Will,” *Journal of Consciousness Studies* 11, no. 7-8 (2004): 162-179. Also, see “Experimental Philosophy,” PhilPapers, ed. James Andow, accessed May 30, 2019, <https://philpapers.org/browse/experimental-philosophy>.

¹⁸⁷ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 27, doi: 10.1136/medethics-2015-102738.

¹⁸⁸ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 27, doi: 10.1136/medethics-2015-102738.

¹⁸⁹ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 27, doi: 10.1136/medethics-2015-102738.

¹⁹⁰ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 27, doi: 10.1136/medethics-2015-102738.

¹⁹¹ Joshua May, “Emotional Reactions to Human Reproductive Cloning,” *Journal of Medical Ethics* 42 (2016): 28, doi: 10.1136/medethics-2015-102738.

¹⁹² James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1130, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.

¹⁹³ James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1131, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.

¹⁹⁴ James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1131, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.

¹⁹⁵ See *x-phi database comment* (endnote number 186) in chapter four for important caveats regarding such claims. For current PhilPapers database, see “Experimental Philosophy,” PhilPapers, ed. James Andow, accessed May 30, 2019, <https://philpapers.org/browse/experimental-philosophy>.

¹⁹⁶ Adalberto De Hoyos, Yareni Monteón, and Myriam M. Altamirano-Bustamante, “Reexamining Healthcare Justice in the Light of Empirical Data,” *Bioethics* 29, no. 9 (2015): 613, accessed January 2, 2017, doi: 10.1111/bioe.12188.

¹⁹⁷ Sigmund Wagner-Tsukamoto, “Consumer Ethics in Japan: An Economic Reconstruction of Moral Agency of Japanese Firms – Qualitative Insights from Grocery/Retail Markets,” *Journal of Business Ethics* 84, no. 1 (2009): 29, accessed January 2, 2017, doi: 10.1007/s10551-008-9671-x.

-
- ¹⁹⁸ Jørgen Weidemann Eriksen, "Should Soldiers Think before They Shoot?" *Journal of Military Ethics* 9, no. 3 (2010): 195, accessed January 2, 2017, doi: 10.1080/15027570.2010.510861.
- ¹⁹⁹ Yixuan Wang, "The Mystery Revealed—Intersectionality in the Black Box: An Analysis of Female Migrants' Employment Opportunities in Urban China," *Hypatia: A Journal of Feminist Philosophy* 30, no. 4 (2015): 862-880, accessed January 2, 2017, doi: 10.1111/hypa.12171.
- ²⁰⁰ Jonathan Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion* (New York: Vintage Books, 2013), 3-31.
- ²⁰¹ Jonathan Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion* (New York: Vintage Books, 2013), 22 and 28.
- ²⁰² Carol Gilligan. *In A Different Voice: Psychological Theory and Women's Development* (Cambridge: Harvard University Press, 2003). ebook.
- ²⁰³ See *x-phi database comment* (endnote number 186) in chapter four for important caveats regarding such claims.
- ²⁰⁴ Sarah von der Mühlen, Tobias Richter, Sebastian Schmid, Elisabeth Marie Schmidt and Kirsten Berthold, "Judging the Plausibility of Arguments in Scientific Texts: A Student–Scientist Comparison," *Thinking & Reasoning* 22, no. 2 (2016): 221-249, accessed June 7, 2017, doi: 10.1080/13546783.2015.1127289.
- ²⁰⁵ Sarah von der Mühlen, Tobias Richter, Sebastian Schmid, Elisabeth Marie Schmidt and Kirsten Berthold, "Judging the Plausibility of Arguments in Scientific Texts: A Student–Scientist Comparison," *Thinking & Reasoning* 22, no. 2 (2016): 221, accessed June 7, 2017, doi: 10.1080/13546783.2015.1127289.
- ²⁰⁶ Sarah von der Mühlen, Tobias Richter, Sebastian Schmid, Elisabeth Marie Schmidt and Kirsten Berthold, "Judging the Plausibility of Arguments in Scientific Texts: A Student–Scientist Comparison," *Thinking & Reasoning* 22, no. 2 (2016): 233, accessed June 7, 2017, doi: 10.1080/13546783.2015.1127289.
- ²⁰⁷ Patricia Barres & P.N. Johnson-Laird, "On Imagining What is True (and What is False)," *Thinking & Reasoning* 9, no. 1 (2003): 1-42, accessed June 7, 2017, doi: 10.1080/13546780244000097.
- ²⁰⁸ Patricia Barres & P.N. Johnson-Laird, "On Imagining What is True (and What is False)," *Thinking & Reasoning* 9, no. 1 (2003): 13, accessed June 7, 2017, doi: 10.1080/13546780244000097.
- ²⁰⁹ Catherine Womack and Norah Mulvaney-Day, "Feminist Bioethics Meets Experimental philosophy: Embracing the Qualitative and Experiential," *The International Journal of Feminist Approaches to Bioethics* 5, no. 1 (2012): 120, accessed December 22, 2016, doi: 10.2979/intjfemappbio.5.1.113.
- ²¹⁰ Catherine Womack and Norah Mulvaney-Day, "Feminist Bioethics Meets Experimental philosophy: Embracing the Qualitative and Experiential," *The International Journal of Feminist Approaches to Bioethics* 5, no. 1 (2012): 120, accessed December 22, 2016, doi: 10.2979/intjfemappbio.5.1.113.
- ²¹¹ Catherine Womack and Norah Mulvaney-Day, "Feminist Bioethics Meets Experimental philosophy: Embracing the Qualitative and Experiential," *The International Journal of Feminist Approaches to Bioethics* 5, no. 1 (2012): 120, accessed December 22, 2016, doi: 10.2979/intjfemappbio.5.1.113.
- ²¹² Catherine Womack and Norah Mulvaney-Day, "Feminist Bioethics Meets Experimental philosophy: Embracing the Qualitative and Experiential," *The International Journal of Feminist Approaches to Bioethics* 5, no. 1 (2012): 120, accessed December 22, 2016, doi: 10.2979/intjfemappbio.5.1.113.
- ²¹³ Elizabeth O'Neill and Edouard Machery, "Introduction: Experimental Philosophy: What is it Good For?" in *Current Controversies in Experimental Philosophy*, ed. Elizabeth O'Neill and Edouard Machery (New York: Routledge, 2014), xxii.
- ²¹⁴ Eddy Nahmias et al., "Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility," *Philosophical Psychology* 18, no. 5 (2005): 572, doi: 10.1080/09515080500264180.
- ²¹⁵ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 172.
- ²¹⁶ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 176.
- ²¹⁷ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 175-176.
- ²¹⁸ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 176.
- ²¹⁹ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 175-177.
- ²²⁰ Eddy Nahmias et al., "The Phenomenology of Free Will," *Journal of Consciousness Studies* 11, no. 7-8 (2004): 172.
- ²²¹ James Andow, "Qualitative Tools and Experimental Philosophy," *Philosophical Psychology* 29, no. 8 (2016): 1131, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.

-
- ²²² James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1129, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.
- ²²³ James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1130, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.
- ²²⁴ See *x-phi database comment* (endnote number 186) in chapter four for important caveats regarding such claims.
- ²²⁵ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), xi. Andow refers to this book by Van Someren et al. in particular.
- ²²⁶ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 1-2.
- ²²⁷ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 1-4.
- ²²⁸ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 4-6.
- ²²⁹ Again, the idea to apply the think aloud method to experimental philosophy is credited to James Andow. For specific details, see James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1130, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.
- ²³⁰ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 29.
- ²³¹ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 20.
- ²³² Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 20.
- ²³³ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 29.
- ²³⁴ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 30.
- ²³⁵ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 30.
- ²³⁶ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 30.
- ²³⁷ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 30.
- ²³⁸ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 4.
- ²³⁹ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 2.
- ²⁴⁰ Philippa Foot “The Problem of Abortion and the Doctrine of the Double Effect,” *Oxford Review* 5 (1967): 5-15, accessed June 5, 2017.
- ²⁴¹ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 37.
- ²⁴² Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 26.
- ²⁴³ “Experimental Philosophy,” PhilPapers, ed. Wesley Buckwalter, accessed January 2, 2017, <https://philpapers.org/browse/experimental-philosophy>. Another example of this structure of definition for experimental philosophy comes from Joshua Alexander’s book *Experimental Philosophy: An Introduction*. In the following excerpt, note that the first sentence paints experimental philosophy broadly and abstractly, while the second sentence adds a more concrete description to the definition: “In recent years, *experimental philosophy* has emerged as an exciting new approach to the study of people’s philosophical intuitions. Experimental philosophers apply the methods of the social and cognitive sciences to the study of philosophical cognition since these methods are better suited than introspective methods to the study of what people, especially *other* people, actually think”; Joshua Alexander, *Experimental Philosophy: An Introduction* (Cambridge: Polity Press, 2012), 2.
- ²⁴⁴ “Experimental Philosophy,” PhilPapers, ed. Wesley Buckwalter, accessed January 2, 2017, <https://philpapers.org/browse/experimental-philosophy>.
- ²⁴⁵ “Experimental Philosophy,” PhilPapers, ed. Wesley Buckwalter, accessed January 2, 2017, <https://philpapers.org/browse/experimental-philosophy>.

-
- ²⁴⁶ Joshua Knobe and Shaun Nichols, “An Experimental Philosophy Manifesto,” in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 3.
- ²⁴⁷ Joshua Knobe and Shaun Nichols, “An Experimental Philosophy Manifesto,” in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 3.
- ²⁴⁸ Joshua Knobe and Shaun Nichols, “An Experimental Philosophy Manifesto,” in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 14.
- ²⁴⁹ James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1131, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.
- ²⁵⁰ James Andow, “Qualitative Tools and Experimental Philosophy,” *Philosophical Psychology* 29, no. 8 (2016): 1131, accessed December 22, 2016, doi: 10.1080/09515089.2016.1224826.
- ²⁵¹ Maarten W. Van Someren, Yvonne F. Barnard, and Jacobijn A. C. Sandberg, *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), xi.
- ²⁵² I based some of my interview process on information from Steinar Kvale’s book *InterViews*. Kvale describes the interview as follows: “The qualitative research interview attempts to understand the world from the subjects’ points of view, to unfold the meaning of peoples’ experiences, to uncover their lived world prior to scientific explanations. [paragraph break] The qualitative research interview is a construction site of knowledge. An interview is literally an *inter view*, an inter change of views between two persons conversing about a theme of mutual interest.” For me, then, the center of the conversation is the experiment and the think aloud session that would have just been performed by the interviewee. For this quote and more information, see Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 1-2.
- ²⁵³ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁵⁴ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁵⁵ Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ²⁵⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. See also: Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ²⁵⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁵⁸ Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ²⁵⁹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁶⁰ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁶¹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁶² Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

²⁶³ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 43.

²⁶⁴ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 43.

²⁶⁵ In both Appendix I and Appendix J there is language indicating that *if* I deem it helpful, I will give participants a second warm-up survey. I ended up giving all participants both warm-up surveys except for one participant because participants truly seemed like they would benefit from it. Perhaps this is unsatisfying, but it is difficult to say precisely how I made the determination for any given participant, so the most I can appeal to is a kind of total assessment of a given situation that I felt confident making while conducting the study. Similarly, I can't say precisely why I asked this or that follow up question in an interview, but I know it was likely based on some assessment of the unique responses a participant produced and how they caught my attention.

²⁶⁶ I have designed this thought experiment to be loosely based on a famous idea often associated with Kant: that of lying being categorically wrong. I also designed it such that it is vaguely similar to the "ought" implies "can" vignettes from the Chituc et al. study. Lastly, I made sure to create a thought experiment that avoids any overly abstract behaviors or anything that might provoke a strong emotional reaction in subjects. For Kant's brief writing on lying, see Immanuel Kant, "On a Supposed Right to Lie from Altruistic Motives," 1-4 accessed via PhilPapers June 6, 2017, <http://www.mc.maricopa.edu/~davpy35701/text/kant-sup-right-to-lie.pdf>.

²⁶⁷ Supplementary material for experiment 3, from Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

²⁶⁸ Supplementary material for experiment 3, from Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

²⁶⁹ Supplementary material for experiment 3, from Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

²⁷⁰ My study; participant 12's interview.

²⁷¹ In designing this follow-up survey, I drew ideas from a few sources I want to acknowledge. Specifically, in designing the note in red font at the top of the survey (which tells participants they can leave responses blank if they prefer) and the last three questions, I drew ideas and wording from the web resource *PsyToolkit*. Specifically, I incorporated this note about leaving questions blank, the specific wording for questions 5 and 6, and the general wording for question 7 after reading this site's guide on measuring demographics. (See "Demographics," *PsyToolkit*, accessed February 2, 2018, <https://www.psychtoolkit.org/survey-library/demographics.html>.) Also, in designing the first four questions on this survey, I drew ideas from a lecture by Zach Weiss on experimental philosophy. Specifically, I incorporated the idea of asking participants how many philosophy classes they have taken, their education level, and some kind of scale-based question of their familiarity with philosophical concepts (and the debate regarding "ought implies can" in particular for my study). (See Zach Weiss, Lecturer, *A Survey of Experimental Philosophy and Epistemology*, YouTube, 15:21, April 26, 2015, <https://www.youtube.com/watch?v=YMDaoudmoEU>.)

²⁷² My study; participant 12's interview.

²⁷³ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 24, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Compare my results to data on this page.

²⁷⁴ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

²⁷⁵ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

²⁷⁶ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

- ²⁷⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁷⁸ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁷⁹ “Easy Fisher Exact Test Calculator,” *Social Science Statistics*, accessed March 11, 2019, <https://www.socscistatistics.com/tests/fisher/default2.aspx>.
- ²⁸⁰ “Effect Size,” Emory Psychology, accessed May 8, 2019, <http://www.psychology.emory.edu/clinical/blwise/Tutorials/SCATTER/scatterplots/effect.htm>.
- ²⁸¹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁸² To reiterate what I report in endnote 266, I designed this thought experiment to be loosely based on a famous idea often associated with Kant: that of lying being categorically wrong. I also designed it such that it is vaguely similar to the “ought” implies “can” vignettes from the Chituc et al. study. For Kant’s brief writing on lying, see Immanuel Kant, “On a Supposed Right to Lie from Altruistic Motives,” 1-4 accessed via PhilPapers June 6, 2017, <http://www.mc.maricopa.edu/~davpy35701/text/kant-sup-right-to-lie.pdf>.
- ²⁸³ To reiterate, this survey is, except for a few small modifications detailed above in this chapter, borrowed from the moral obligation/high fault/unable survey from Chituc et al.’s third experiment. For the original survey details and more information, see Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ²⁸⁴ To reiterate, this survey is, except for a few small modifications detailed above in this chapter, borrowed from the moral obligation/low fault/unable survey from Chituc et al.’s third experiment. For the original survey details and more information, see Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ²⁸⁵ To reiterate what is in endnote number 271, I drew key features of my follow-up survey from two resources: the web resource *PsyToolkit* (see “Demographics,” *PsyToolkit*, accessed February 2, 2018, <https://www.psychology.emory.edu/clinical/blwise/Tutorials/SCATTER/scatterplots/effect.htm>) and also from a lecture by Zach Weiss (see Zach Weiss, Lecturer, *A Survey of Experimental Philosophy and Epistemology*, YouTube, 15:21, April 26, 2015, <https://www.youtube.com/watch?v=YMDaoudmoEU>). Please see endnote number 271 for further information on how I derived features of my survey from these sources.
- ²⁸⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁸⁷ See Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 249-261; and Walter Sinnott-Armstrong, “‘Ought to Have’ and ‘Could Have,’” *Analysis* 45, no. 1 (1985): 44-48.
- ²⁸⁸ Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *The Philosophical Review* 93, no. 2, (1984): 249-261.
- ²⁸⁹ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁹⁰ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ²⁹¹ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, “‘Ought Implies Can,’ Framing Effects, and ‘Empirical Refutations,’” forthcoming in *Philosophia* 46, no. 1 (2018): 165-182. See Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.; See also Paul Henne et al., “An Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Analysis* 76, no. 3 (2016): 283-290.

-
- ²⁹² Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.; See also Paul Henne et al., “An Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Analysis* 76, no. 3 (2016): 283-290.
- ²⁹³ My study; participant 30’s full transcript.
- ²⁹⁴ My study; participant 30’s think aloud session.
- ²⁹⁵ My study; participant 30’s think aloud session.
- ²⁹⁶ My study; participant 30’s interview.
- ²⁹⁷ My study; participant 30’s think aloud session.
- ²⁹⁸ My study; participant 30’s think aloud session.
- ²⁹⁹ In one work, Sinnott-Armstrong presents an argument against OIC wherein “implies” is semantic. In the course of this argument he deals with some interesting features of temporality and language. There might be more to explore here in a larger project that connects x-phi research back to the larger debate as to whether OIC is the case, but that isn’t my present concern. Also, my reason for focusing on temporality is different than is Sinnott-Armstrong’s in this paper, as far as I can tell. My focus is simply on getting participants to lock in their judgments of “ought” and “can” at precise times so we aren’t as researchers mislead into thinking they hold a different view than they really do as might be obscured by the trickiness presented in timestamps. See Walter Sinnott-Armstrong, “‘Ought to Have’ and ‘Could Have,’” *Analysis* 45, no. 1 (1985): 44-48.
- ³⁰⁰ Bart Streumer, “I Ought to Reply, So I Can,” *Philosophia* (2018): 6-7, <https://doi.org/10.1007/s11406-018-0042-2>.
- ³⁰¹ Bart Streumer, “I Ought to Reply, So I Can,” *Philosophia* (2018): 7, <https://doi.org/10.1007/s11406-018-0042-2>.
- ³⁰² Leben, Derek, “In Defense of ‘Ought Implies Can,’” in *Oxford Studies in Experimental Philosophy, Volume 2*, eds. Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Oxford: Oxford University Press, 2018), 155, chap. 6, Kindle.
- ³⁰³ My study; participant 30’s interview.
- ³⁰⁴ My study; participant 14’s full transcript.
- ³⁰⁵ My study; participant 14’s think aloud session.
- ³⁰⁶ My study; participant 14’s interview.
- ³⁰⁷ My study; participant 14’s think aloud session.
- ³⁰⁸ My study; participant 14’s interview.
- ³⁰⁹ My study; participant 14’s think aloud session.
- ³¹⁰ My study; participant 14’s think aloud session.
- ³¹¹ My study; participant 14’s interview.
- ³¹² My study; participant 14’s interview.
- ³¹³ My study; participant 14’s interview.
- ³¹⁴ My study; participant 4’s interview.
- ³¹⁵ My study; participant 5’s interview.
- ³¹⁶ My study; participant 13’s interview.
- ³¹⁷ My study; participant 13’s interview.
- ³¹⁸ My study; participant 13’s interview.
- ³¹⁹ My study; participant 13’s interview.
- ³²⁰ My study; participant 13’s interview.
- ³²¹ My study; participant 13’s interview.
- ³²² My study; participant 13’s interview.
- ³²³ Joshua Knobe and Shaun Nichols, “An Experimental Philosophy Manifesto,” in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 6.
- ³²⁴ My study; participant 22’s interview.
- ³²⁵ My study; participant 7’s interview.
- ³²⁶ My study; participant 23’s think aloud session.
- ³²⁷ My study; participant 23’s interview.
- ³²⁸ My study; participant 23’s interview.
- ³²⁹ Interestingly, when reviewing this transcript, I noticed that participant 2 talks as though he selects “5” for questions 1 and 2, but I have in my records that he selected “2” for each. It is possible that I mis-recorded his quantitative responses, but I think it much more likely that during the interview he talked about his agreement with those statements as hitting on the extremes when in reality he circled a lower level of agreement, that of 2. It is true subjects can look at their surveys as they complete the interviews, but they are also focused on a host of things—i.e.,

talking aloud, tracking the thought experiment, tracking my questions—whereas when I recorded their responses I could do so without thinking of anything else. I made sure to double check all my data entry. In sum, I am going to treat the recorded data as true and accurate because I don't have sufficient reason to think it isn't.

³³⁰ My study; participant 2's interview.

³³¹ My study; participant 2's interview.

³³² My study; participant 10's interview.

³³³ My study; participant 10's interview.

³³⁴ My study; participant 34's interview.

³³⁵ For the specific thought experiment, see Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 249-261.

³³⁶ For more insight on Sinnott-Armstrong's treatment of OIC, see Walter Sinnott-Armstrong, "'Ought to Have' and 'Could Have,'" *Analysis* 45, no. 1 (1985): 44-48.

³³⁷ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

³³⁸ Paul Henne et al., "An Empirical Refutation of 'Ought' Implies 'Can,'" *Analysis* 76, no. 3 (2016): 283-290.

³³⁹ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations,'" forthcoming in *Philosophia* 46, no. 1 (2018): 1-32.

³⁴⁰ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations,'" forthcoming in *Philosophia* 46, no. 1 (2018): 26-27.

³⁴¹ *The Holy Bible: New International Version* (Grand Rapids: Zondervan Bible Publishers, 1996): Exodus 20:1-21.

³⁴² "The Life You Can Save," The Life You Can Save, accessed June 1, 2019, <https://www.thelifeyoucansave.org/>.

³⁴³ Aristotle, *Ethics*, trans. J. A. K. Thomson (London: Penguin Books, 1976).

³⁴⁴ Immanuel Kant, *Fundamental Principles of the Metaphysic of Morals*, trans. Thomas K. Abbott (Indianapolis: The Bobbs-Merrill Company, Inc., 1949).

³⁴⁵ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodé* 4, no. 6 (2015): 248. Though, it is worth noting that, in reply to this reply, Mizrahi provides at least one example of a philosopher, Wedgwood, advocating otherwise. That said, the spirit of this claim still holds it seems. See Mizrahi, Moti, "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith," *Methodé* 4, no. 6 (2015): 252.

³⁴⁶ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodé* 4, no. 6 (2015): 248.

³⁴⁷ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations,'" forthcoming in *Philosophia* 46, no. 1 (2018): 1-32.

³⁴⁸ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations,'" forthcoming in *Philosophia* 46, no. 1 (2018): 13.

³⁴⁹ Moti Mizrahi, "Ought, Can, and Presupposition: An Experimental Study," *Methodé* 4, no. 6 (2015): 232-243.

³⁵⁰ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodé* 4, no. 6 (2015): 245.

³⁵¹ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodé* 4, no. 6 (2015): 247. As previously mentioned in an endnote (number 345), Mizrahi provides one example of such in Wedgwood. Still, Kurthy and Lawford-Smith appear to be on solid ground with the spirit of their claim. See Mizrahi, Moti, "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith," *Methodé* 4, no. 6 (2015): 252.

³⁵² Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodé* 4, no. 6 (2015): 247.

³⁵³ Mizrahi, Moti, "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith," *Methodé* 4, no. 6 (2015): 254.

³⁵⁴ Mizrahi, Moti, "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith," *Methodé* 4, no. 6 (2015): 253.

³⁵⁵ Wesley Buckwalter and John Turri, "Inability and Obligation in Moral Judgment," *PLoS ONE*, 10, no. 8 (2015): 2. doi:10.1371/journal.pone.0136589.

³⁵⁶ Wesley Buckwalter and John Turri, "Inability and Obligation in Moral Judgment," *PLoS ONE*, 10, no. 8 (2015): 3-5. doi:10.1371/journal.pone.0136589.

- ³⁵⁷ Kurthy et al., in a paper I discuss next, make note of something like this: “Aspects of [Buckwalter and Turri’s] design...may make the option ‘obligated, but not able’ the sole plausible answer, though in a way that is not inconsistent with the OIC principle.” See Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 3, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁵⁸ Again, this is something that Kurthy et al.’s work, which I discuss next as a reply to Buckwalter and Turri, speaks to. This participant from their study, which modifies Buckwalter and Turri’s work, who selected that an agent is “‘obligated, but not able’” to carry out an obligation, seems to think the OIC principle upholds so clearly that they almost read as bothered that someone might cast them as possibly negating it: “‘Obviously Walter is no longer obligated to pick up Brown from the airport and anyone who tries to philosophically argue the case is limited in their scope of reality. Walter agreed to pick up someone from the airport but after being severely incapacitated due to a car accident he is no longer able (or obligated) to pick up the person and he should find an alternative.’” See Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 8, no. 4 (2017): 1-24, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁵⁹ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 1-24, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁰ I was happy to read that Kurthy et al. specifically note that their qualitative results help show that participants’ apparently anti-OIC responses aren’t so clearly evidence against OIC. However, I disagree with their wording that it is the participants who misinterpret; I place that burden on the researchers. See Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 21, no. 4 (2017): 1-24, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶¹ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 23, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶² Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 7, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶³ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 8, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁴ My study; participant 30’s interview.
- ³⁶⁵ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 20, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁶ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 11-12, no. 4 (2017): 13, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁷ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 11, no. 4 (2017): 13, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁸ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 8, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁶⁹ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 7, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁷⁰ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, “Does Ought Imply Can?” *PLoS ONE* 12, no. 4 (2017): 13, <https://doi.org/10.1371/journal.pone.0175206>.
- ³⁷¹ Although, potentially problematically, there is such language (as mentioned in endnote 116) in experiment 1 when participants are asked to agree or disagree with the following: “At eleven forty-five, it is still true that Adams ought to meet Brown at noon.” See Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ³⁷² Leben, Derek, “In Defense of ‘Ought Implies Can,’” in *Oxford Studies in Experimental Philosophy, Volume 2*, eds. Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Oxford: Oxford University Press, 2018), 157-159, chap. 6, Kindle.
- ³⁷³ Leben, Derek, “In Defense of ‘Ought Implies Can,’” in *Oxford Studies in Experimental Philosophy, Volume 2*, eds. Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Oxford: Oxford University Press, 2018), 160-163, chap. 6, Kindle.
- ³⁷⁴ Leben, Derek, “In Defense of ‘Ought Implies Can,’” in *Oxford Studies in Experimental Philosophy, Volume 2*, eds. Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Oxford: Oxford University Press, 2018), 160-163, chap. 6, Kindle.
- ³⁷⁵ I just want to give some brief appreciation for Hannon who, in the end, offers a measured analysis of experimental philosophy and points to some creative and more expansive methodologies that x-philosophers might use, some of which get at my support for methods that are more probing than the basic quantitative survey. Michael

Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4147-4168, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁷⁶ Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4150-4154, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁷⁷ Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4152-4153, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁷⁸ As I mentioned in previous endnotes (numbers 116 and 371), this language of "still" isn't ideal as it might bias a kind of interpretation of obligation as a thing that goes away. See Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013; also see Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4153, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁷⁹ Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4154, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁸⁰ Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4153, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁸¹ Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4153, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁸² Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4153, <https://doi.org/10.1007/s11229-017-1412-1>.

³⁸³ My study; participant 30's think aloud session.

³⁸⁴ Yishai Cohen, "An Analysis of Recent Empirical Data on 'Ought' Implies 'Can,'" *Philosophia* (2017): 1-11. DOI 10.1007/s11406-017-9892-2.

³⁸⁵ I've brought up the significance of such subtle language throughout this project, and I refer to this instance in particular in previous endnotes (numbers 116, 371, and 378).

³⁸⁶ Yishai Cohen, "An Analysis of Recent Empirical Data on 'Ought' Implies 'Can,'" *Philosophia* (2017): 6. DOI 10.1007/s11406-017-9892-2.

³⁸⁷ Tamar Szabó Gendler and Shen-yi Liao, "The Problem of Imaginative Resistance," in *The Routledge Companion to Philosophy of Literature*, eds. Noël Carroll and John Gibson (New York: Routledge, 2015): 405, doi-org.ccl.idm.oclc.org/10.4324/9781315708935.

³⁸⁸ As far as I can tell, there is a straightforward connection between the examples I explore and the phenomenon of imaginative resistance as just defined, but I must note that much of the discussion around the topic, including in the chapter from where this definition comes, focuses on resistance to morally objectionable or morally impossible fiction, such as fiction describing the permissibility of female infanticide. It seems to me the focus on morally objectionable fiction is not intended to constrain the phenomenon to only such fictions, but I want to acknowledge that the surrounding discussion of imaginative resistance I've encountered is coupled with morally objectionable fiction and not with the kinds of thought experiment constraints I explore. As just another example of this, the summary of the PhilPapers section dedicated to imaginative resistance gives an example regarding the moral rightness of female infanticide. See "Imaginative Resistance," PhilPapers, ed. Amy Kind, accessed April 8, 2019, <https://philpapers.org/browse/imaginative-resistance>.

³⁸⁹ Tamar Szabó Gendler and Shen-yi Liao, "The Problem of Imaginative Resistance," in *The Routledge Companion to Philosophy of Literature*, eds. Noël Carroll and John Gibson (New York: Routledge, 2015): 409, doi-org.ccl.idm.oclc.org/10.4324/9781315708935. In discussing such theories, Gendler and Liao quickly move to examples wherein someone is asked to imagine something morally objectionable, namely the permissibility of female infanticide. Again, this doesn't, on its face, suggest that imaginative resistance doesn't apply when thinking about participants resisting the time constraints of a thought experiment, but it is worth highlighting that canonical examples of imaginative resistance appear to be focused on morally objectionable fictions.

³⁹⁰ Tamar Szabó Gendler and Shen-yi Liao, "The Problem of Imaginative Resistance," in *The Routledge Companion to Philosophy of Literature*, eds. Noël Carroll and John Gibson (New York: Routledge, 2015): 407, doi-org.ccl.idm.oclc.org/10.4324/9781315708935.

³⁹¹ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

³⁹² Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

³⁹³ From participant 10. See Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

³⁹⁴ From participant 184. See Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

³⁹⁵ From participant 195. See Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

³⁹⁶ Full vignettes for OIC experiment 3 are available as supplementary material to the study; see Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

³⁹⁷ My study; participant 18’s think aloud session.

³⁹⁸ My study; participant 19’s think aloud session.

³⁹⁹ My study; participant 20’s think aloud session.

⁴⁰⁰ The reader may notice that I’ve not focused on the quantitative responses that correlate with these qualitative responses. This is because I’m interested in seeing hints of the influence of imaginative resistance whether or not it appears reflected in the quantitative responses. If a participant flirts with negating the constraints of the vignette but selects that Brown can’t make it, it should still be a worry as to whether his or her judgment was one of “can’t”—the participant might be selecting the “correct” response of “can’t” all the while harboring some internal reasoning that perhaps Brown *can* make it.

⁴⁰¹ Shaun Nichols and Joshua Knobe, “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions,” in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 115-116.

⁴⁰² Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁰³ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁰⁴ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁰⁵ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁰⁶ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁰⁷ From participant 23. See Supplementary material for experiment 1, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

⁴⁰⁸ From participant 8. See Supplementary material for experiment 1, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

-
- ⁴⁰⁹ Although, I might note that philosophers too have issues with accepting that impossible means impossible. The most compelling argument I'm aware of in support of OIC essentially points out that if those who oppose OIC truly appreciated what it means to say that an obligation is impossible to fulfil, then they'd accept OIC (or at least have one fewer reason to reject it). In reference to Chituc et al.'s OIC experiment 1, Streumer presents an argument that essentially points out that just as one would conclude that an agent no longer "ought" to fulfil an obligation to meet at time *t* once time *t* has passed, since to do so would be to obligate agent to change the past, so too is it the case that agent no longer "ought" to fulfil an obligation at the first moment prior to time *t* in which agent no longer has time to get to the meeting by time *t* since that too would require agent to change the past, thereby making themselves available to leave at or before the last moment in which he could still make the meeting. In short, there is reason to think that some philosophers might actually switch to defenders of OIC were they to properly appreciate that impossibility (i.e., of changing the past) is just that. Changing a past deadline is as impossible—that is to say it is impossible—as changing past events so as to make a future deadline meet-able. For more details, see Bart Streumer, "I Ought to Reply, So I Can," *Philosophia* (2018): 6-7, <https://doi.org/10.1007/s11406-018-0042-2>.
- ⁴¹⁰ My study; participant 12's interview.
- ⁴¹¹ My study; participant 18's interview.
- ⁴¹² My study; participant 34's interview.
- ⁴¹³ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 21-22, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ⁴¹⁴ This is an allusion to the trolley problem. For more details see Philippa Foot "The Problem of Abortion and the Doctrine of the Double Effect," *Oxford Review* 5 (1967): 5-15, accessed June 5, 2017.
- ⁴¹⁵ I based this case on the trolley problem. For more details see Philippa Foot "The Problem of Abortion and the Doctrine of the Double Effect," *Oxford Review* 5 (1967): 5-15, accessed June 5, 2017.
- ⁴¹⁶ Robert Kirk, "Zombies," *The Stanford Encyclopedia of Philosophy*, Spring 2019, <https://plato.stanford.edu/entries/zombies/>.
- ⁴¹⁷ Alicia Kissinger-Knox, Patrick Aragon, and Moti Mizrahi, "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations,'" forthcoming in *Philosophia* 46, no. 1 (2018): 6-7.
- ⁴¹⁸ Shaun Nichols and Joshua Knobe, "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions," in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 111.
- ⁴¹⁹ Shaun Nichols and Joshua Knobe, "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions," in *Experimental Philosophy*, eds. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 114-115.
- ⁴²⁰ Walter Sinnott-Armstrong, "Abstract + Concrete = Paradox," in *Experimental Philosophy*, ed. Joshua Knobe and Shaun Nichols (New York: Oxford University Press, 2008), 209-229.
- ⁴²¹ Tehila Kogut and Ilana Ritov, "The Singularity Effect of Identified Victims in Separate and Joint Evaluations," *Organizational Behavior and Human Decision Processes* 97 (2005): 106-116, doi: 10.1016/j.obhdp.2005.02.003.
- ⁴²² My study; participant 33's full interview.
- ⁴²³ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252-254.
- ⁴²⁴ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.
- ⁴²⁵ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252.
- ⁴²⁶ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252.
- ⁴²⁷ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252.
- ⁴²⁸ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252.
- ⁴²⁹ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 252-253.
- ⁴³⁰ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 258.

-
- ⁴³¹ My study; participant 33's interview.
- ⁴³² My study; participant 33's interview.
- ⁴³³ My study; participant 33's interview.
- ⁴³⁴ Some of the participant responses from Kurthy et al.'s work supports the intuition/judgment that participants view obligation up until the moment of inability. However, as I lamented before, the language in their studies might be biasing this interpretation (at least in some cases) in a way that Chituc et al.'s questions, and my follow-up question about 5:55, does not. For example, I worry about the language of "still" such that participants who are asked if they think an agent "still" ought to do something or be viewed as under an obligation might take this as invitation to divide up a scenario into before and after the inability. At the same time, they do report one qualitative response in connection with a scenario that seems to align clearly with Sinnott-Armstrong's intuition. The reader is given the impression that this is a rare occurrence, and it is the only highlighted qualitative response that clearly echoes the negation of OIC. In the scenario, an agent is to save a drowning girl before succumbing to paralysis; one participant says, "Well I assume nothing has changed about the girls [sic] situation just because Michael can't move so the obligation to save her is still there, even if he can't move it still exists." See Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 1-24, no. 4 (2017): 1-24, page 18 for the quoted participant, <https://doi.org/10.1371/journal.pone.0175206>.
- ⁴³⁵ My study; participant 14's think aloud session.
- ⁴³⁶ My study; participant 14's think aloud session.
- ⁴³⁷ Leben, Derek, "In Defense of 'Ought Implies Can,'" in *Oxford Studies in Experimental Philosophy, Volume 2*, eds. Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Oxford: Oxford University Press, 2018), 157-159, chap. 6, Kindle.
- ⁴³⁸ My study; participant 31's interview.
- ⁴³⁹ My study; participant 15's interview.
- ⁴⁴⁰ My study; participant 2's interview.
- ⁴⁴¹ My study; participant 6's interview.
- ⁴⁴² Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 261.
- ⁴⁴³ Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2 (1984): 261.
- ⁴⁴⁴ My study; participant 6's interview.
- ⁴⁴⁵ My study; participant 6's interview.
- ⁴⁴⁶ My study; participant 6's interview.
- ⁴⁴⁷ My study; participant 14's interview.
- ⁴⁴⁸ My study; participant 14's interview.
- ⁴⁴⁹ My study; participant 11's interview.
- ⁴⁵⁰ My study; participant 28's interview.
- ⁴⁵¹ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 11, no. 4 (2017): 11-12, <https://doi.org/10.1371/journal.pone.0175206>.
- ⁴⁵² Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 11, no. 4 (2017): 12-14, <https://doi.org/10.1371/journal.pone.0175206>.
- ⁴⁵³ My study; participant 38's interview.
- ⁴⁵⁴ My study; participant 8's think aloud session.
- ⁴⁵⁵ One added value in Buckwalter and Turri's experimentation and also Kurthy et al.'s is the use of scenarios where more serious consequences are on the line, such as the life of drowning swimmers. It would be interesting to see how the think aloud method and interview might reveal different insights into people's OIC judgments were they applied to such cases; a movie date really just isn't that serious. See Wesley Buckwalter and John Turri, "Inability and Obligation in Moral Judgment," *PLoS ONE*, 10, no. 8 (2015): 1-20. doi:10.1371/journal.pone.0136589; and see also Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 12, no. 4 (2017): 1-24, <https://doi.org/10.1371/journal.pone.0175206>.
- ⁴⁵⁶ My study; participant 12's think aloud session.
- ⁴⁵⁷ My study; participant 36's think aloud session.
- ⁴⁵⁸ My study; participant 28's interview.
- ⁴⁵⁹ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodes* 4, no. 6 (2015): 245.
- ⁴⁶⁰ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodes* 4, no. 6 (2015): 245.

-
- ⁴⁶¹ Miklos Kurthy and Holly Lawford-Smith, "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study,'" *Methodes* 4, no. 6 (2015): 245.
- ⁴⁶² My study; participant 33's interview.
- ⁴⁶³ My study; participant 33's interview.
- ⁴⁶⁴ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 7, doi: 10.1007/s11098-018-1159-7.
- ⁴⁶⁵ Buckwalter mentions a few examples of indirect control agents at institutions can take when considering implicit biases: "Anonymous grading, employee evaluation, or resume searchers, for instance, could prevent biases from influencing evaluations." From Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 7, doi: 10.1007/s11098-018-1159-7.
- ⁴⁶⁶ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 13-14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁶⁷ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 13, doi: 10.1007/s11098-018-1159-7.
- ⁴⁶⁸ See Walter Sinnott-Armstrong, "'Ought' Conversationally Implies 'Can,'" *The Philosophical Review* 93, no. 2, (1984): 249-261.
- ⁴⁶⁹ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 13-14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁷⁰ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁷¹ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁷² Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁷³ Wesley Buckwalter, "Implicit Attitudes and the Ability Argument," *Philosophical Studies*, September (2018): 14, doi: 10.1007/s11098-018-1159-7.
- ⁴⁷⁴ Supplementary material for experiment 3, from Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- ⁴⁷⁵ My study; participant 30's interview. In her think aloud session, she specifically attributes blame to Brown for making a conscious decision: "Brown is to blame for not making it to the theater by 6. Yes 'cause he consciously made the decision not to get there on time."
- ⁴⁷⁶ My study; participant 7's think aloud session.
- ⁴⁷⁷ My study; participant 3's think aloud session.
- ⁴⁷⁸ My study; participant 8's think aloud session.
- ⁴⁷⁹ My study; participant 7's think aloud session.
- ⁴⁸⁰ My study; participant 5's interview.
- ⁴⁸¹ My study; participant 15's interview.
- ⁴⁸² Michael Hannon, "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (2018): 4150-4154, <https://doi.org/10.1007/s11229-017-1412-1>.
- ⁴⁸³ Rosalind Hursthouse, "Fallacies and Moral Dilemmas," *Argumentation* 9, no. 4 (1995): 619.
- ⁴⁸⁴ Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa, "Does Ought Imply Can?" *PLoS ONE* 11, no. 4 (2017): 11-12, <https://doi.org/10.1371/journal.pone.0175206>.
- ⁴⁸⁵ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Chituc et al. reference here the work of Alicke and others. For more information on blame validation, see Mark D. Alicke, "Culpable Control and the Psychology of Blame," *Psychological Bulletin* 126, no. 4 (2000): 556-574. doi: 10.1037/W0033-2909.126.4.556.
- ⁴⁸⁶ Vladimir Chituc et al., "Blame, Not Ability, Impacts Moral 'Ought' Judgments for Impossible Actions: Toward an Empirical Refutation of 'Ought' Implies 'Can,'" *Cognition* 150 (2016): 21, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Chituc et al. here refer to the work by Turri and Blouw on excuse validation. See John Turri and Peter Blouw, "Excuse Validation: A Study in Rule-breaking," *Philosophical Studies* 172, no. 3 (2015): 615-634. doi: 10.1007/s11098-014-0322-z.

⁴⁸⁷ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁸⁸ Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 23-24, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013.

⁴⁸⁹ My study; participant 8’s interview.

⁴⁹⁰ My study; participant 8’s interview.

⁴⁹¹ My study; participant 8’s think aloud session.

⁴⁹² Paulette M. Rothbauer, “Triangulation,” in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 1-4. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.

⁴⁹³ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 33.

⁴⁹⁴ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 33.

⁴⁹⁵ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 33.

⁴⁹⁶ Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 33.

⁴⁹⁷ I got the idea to refer to this as the “logistical ought” from reading participant 36’s transcript, which I will refer to just below in discussion of multiple meanings of “ought.”

⁴⁹⁸ I want to thank Brian Keeley for reminding me that some people, perhaps those from the American South or heavily steeped in its culture or dialects, might understand the word “ought” as having these two meanings such that the OIC literature, including my study, might be generally missing an important way that some people think about “ought.” Keeley’s comments inspired this whole discussion of the logistical ought here in my concluding chapter. In addition, I want to note that there is likely much more complexity to this logistical ought than I am unpacking here, and future x-phi research on OIC would benefit from conceptually analyzing it, or drawing from established analysis of it, in service of experimental studies. For now, though, I am only concerned with providing a rough outline of the logistical ought in order to discuss the limitations of my current project.

⁴⁹⁹ Thanks to Brian Keeley for bringing up to me the possibility of using a different verb (or verb phrase) with this question in particular.

⁵⁰⁰ Wesley Buckwalter and John Turri, “Inability and Obligation in Moral Judgment,” *PLoS ONE*, 10, no. 8 (2015): 1-20. doi:10.1371/journal.pone.0136589.

⁵⁰¹ My study; participant 36’s think aloud session.

⁵⁰² My study; participant 36’s think aloud session.

⁵⁰³ My study; participant 30’s interview.

⁵⁰⁴ To emphasize what I report in endnote 266, I designed this thought experiment to be loosely based on a famous idea often associated with Kant: that of lying being categorically wrong. I also designed it such that it is vaguely similar to the “ought” implies “can” vignettes from the Chituc et al. study. For Kant’s brief writing on lying, see Immanuel Kant, “On a Supposed Right to Lie from Altruistic Motives,” 1-4 accessed via PhilPapers June 6, 2017, <http://www.mc.maricopa.edu/~davpy35701/text/kant-sup-right-to-lie.pdf>.

⁵⁰⁵ This survey is, except for a few small modifications detailed in chapter five, borrowed from the moral obligation/high fault/unable survey from Chituc et al.’s third experiment. For the original survey details and more information, see Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

⁵⁰⁶ This survey is, except for a few small modifications detailed in chapter five, borrowed from the moral obligation/low fault/unable survey from Chituc et al.’s third experiment. For the original survey details and more information, see Supplementary material for experiment 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.

⁵⁰⁷ To emphasize what is in endnote number 271, I drew key features of my follow-up survey from two resources: the web resource *PsyToolkit* (see “Demographics,” *PsyToolkit*, accessed February 2, 2018,

<https://www.psytoolkit.org/survey-library/demographics.html>) and also from a lecture by Zach Weiss (see Zach Weiss, Lecturer, *A Survey of Experimental Philosophy and Epistemology*, YouTube, 15:21, April 26, 2015, <https://www.youtube.com/watch?v=YMDaoudmoEU>). Please see endnote number 271 for further information on how I derived features of my survey from these sources.

⁵⁰⁸ Some of my thinking about and choices regarding transcripts emerged from reading the following sources: Michelle K. McGinn, “Transcript,” in *The SAGE Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, Inc., 2008), 1-4. <http://dx.doi.org/10.4135/9781412963909.n462>.; Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994).; Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996).

⁵⁰⁹ Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 160.

⁵¹⁰ Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 165.

⁵¹¹ Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 166.

⁵¹² Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 166.

⁵¹³ Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 167-168.

⁵¹⁴ Steinar Kvale, *InterViews: An Introduction to Qualitative Research Interviewing* (Thousand Oaks: Sage Publications, 1996), 169.

⁵¹⁵ Because of Van Someren’s suggestion on page 45, I opted for “unintelligible” as the preferred term for things unintelligible. See Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 45.

⁵¹⁶ My attempt to use little punctuation aligns heavily with Van Someren’s guidelines on page 46. See Van Someren et al., *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes* (London: Academic Press, 1994), 46.

⁵¹⁷ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See tab with data minus failed checks.

⁵¹⁸ Supplementary material for experiment 2, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can,’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>. See tab with data minus failed checks.

References

- Aristotle. *Ethics*. Translated by J. A. K. Thomson. London: Penguin Books, 1976.
- “about experimental philosophy.” *experimental philosophy*. Accessed June 2, 2017, <http://experimental-philosophy.yale.edu/ExperimentalPhilosophy.html>.
- Alexander, Joshua. *Experimental Philosophy: An Introduction*. Cambridge: Polity Press, 2012.
- Alicke, Mark D. “Culpable Control and the Psychology of Blame.” *Psychological Bulletin* 126, no. 4 (2000): 556-574. doi: 10.1037W 0033-2909.126.4.556.
- Andow, James. “Qualitative Tools and Experimental Philosophy.” *Philosophical Psychology* 29, no. 8 (2016): 1128-1141. Accessed December 22, 2016. doi: 10.1080/09515089.2016.1224826.
- Andow, James, and Florian Cova. “Why Compatibilist Intuitions Are Not Mistaken: A Reply to Feltz and Millan.” *Philosophical Psychology* 29, no. 4 (2016): 550-566. doi: 10.1080/09515089.2015.1082542.
- Appiah, Kwame Anthony. “The New New Philosophy.” Idea Lab. *The New York Times Magazine*, December 9, 2007. <https://www.nytimes.com/2007/12/09/magazine/09wwln-idealab-t.html>.
- Barres, Patricia and P.N. Johnson-Laird. “On Imagining What is True (and What is False).” *Thinking & Reasoning* 9, no. 1 (2003): 1-42. Accessed June 7, 2017, doi: 10.1080/13546780244000097.
- Buckwalter, Wesley. “Implicit Attitudes and the Ability Argument.” *Philosophical Studies*, September (2018): 1-30. doi: 10.1007/s11098-018-1159-7.
- Buckwalter, Wesley and John Turri. “Inability and Obligation in Moral Judgment.” *PLoS ONE*, 10, no. 8 (2015): 1-20. doi:10.1371/journal.pone.0136589.
- Chituc, Vlad and Paul Henne. “The Data Against Kant.” *New York Times*, February 19, 2016.
- Chituc, Vladimir, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard. “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can.’” *Cognition* 150 (2016): 20-25. Accessed June 5, 2017. doi: 10.1016/j.cognition.2016.01.013.
- Cohen, Yishai. “An Analysis of Recent Empirical Data on ‘Ought’ Implies ‘Can.’” *Philosophia* (2017): 1-11. DOI 10.1007/s11406-017-9892-2.

- Cokely, Edward T., and Colleen M. Kelley. "Cognitive Abilities and Superior Decision Making Under Risk: A Protocol Analysis and Process Model Evaluation." *Judgment and Decision Making* 4, no. 1 (2009): 20-33.
- Creswell, John W. "Mixed Methods Research." In *The SAGE Encyclopedia of Qualitative Research Methods*, edited by Lisa M. Given, 1-5, Thousand Oaks: SAGE Publications, Inc., 2008. <http://dx.doi.org/10.4135/9781412963909.n269>.
- De Hoyos, Adalberto, Yareni Monteón and Myriam M. Altamirano-Bustamante. "Reexamining Healthcare Justice in the Light of Empirical Data." *Bioethics* 29, no. 9 (2015): 613-621. Accessed January 2, 2017. doi: 10.1111/bioe.12188.
- "Demographics," *PsyToolkit*. Accessed February 2, 2018. <https://www.psychology.emory.edu/clinical/bliwise/Tutorials/SCATTER/scatterplots/effect.htm>.
- "Denmark: Percent Christian." *The Global Economy*. Accessed April 24, 2019. <https://www.theglobaleconomy.com/Denmark/christians/>.
- "Easy Fisher Exact Test Calculator." *Social Science Statistics*. Accessed March 11, 2019. <https://www.socscistatistics.com/tests/fisher/default2.aspx>.
- "Effect Size." Emory Psychology. Accessed May 8, 2019. <http://www.psychology.emory.edu/clinical/bliwise/Tutorials/SCATTER/scatterplots/effect.htm>.
- Eriksen, Jørgen Weidemann. "Should Soldiers Think before They Shoot?" *Journal of Military Ethics* 9, no. 3 (2010): 195-218. Accessed January 2, 2017. doi: 10.1080/15027570.2010.510861.
- Feltz, Adam, and Melissa Millan. "An Error Theory for Compatibilist Intuitions." *Philosophical Psychology* 28, no. 4 (2015): 529-555. doi: 10.1080/09515089.2013.865513.
- Foot, Philippa. "The Problem of Abortion and the Doctrine of the Double Effect" *Oxford Review* 5 (1967): 5-15. Accessed June 5, 2017.
- Gendler, Tamar Szabó, and Shen-yi Liao. "The Problem of Imaginative Resistance." In *The Routledge Companion to Philosophy of Literature*, edited by Noël Carroll and John Gibson, 405-418. New York: Routledge, 2015. doi-org.ccl.idm.oclc.org/10.4324/9781315708935.
- Gettier, Edmund L. "Is Justified True Belief Knowledge?" *Analysis* 23, no. 6 (1963): 121-123.
- Gilligan, Carol. *In A Different Voice: Psychological Theory and Women's Development*. Cambridge: Harvard University Press, 2003. ebook.

- Goldman, Alvin I. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73, no. 20 (1976): 771-791.
- Haidt, Jonathan. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Vintage Books, 2013.
- Hannon, Michael. "Intuitions, Reflective Judgments, and Experimental Philosophy." *Synthese* 195 (2018): 4147-4168. <https://doi.org/10.1007/s11229-017-1412-1>.
- Henne, Paul, Vladimir Chituc, Felipe De Brigard, and Walter Sinnott-Armstrong. "An Empirical Refutation of 'Ought' Implies 'Can.'" *Analysis* 76, no. 3 (2016): 283-290.
- The Holy Bible: New International Version*. Grand Rapids: Zondervan Bible Publishers, 1996.
- Hughes, Rhidian. "Vignettes." In *The SAGE Encyclopedia of Qualitative Research Methods*, edited by Lisa M. Given, 1-5. Thousand Oaks: SAGE Publications, Inc., 2008. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n482>.
- Hursthouse, Rosalind. "Fallacies and Moral Dilemmas." *Argumentation* 9, no. 4 (1995): 617-632.
- Inglehart, Ronald, Miguel Basañez, and Alejandro Moreno. *Human Values and Beliefs: A Cross-Cultural Sourcebook*. Ann Arbor: University of Michigan Press, 1998. ebook.
- Kant, Immanuel. *Critique of Pure Reason*. Translated by Norman Kemp Smith. New York: Palgrave Macmillan, 2007.
- Kant, Immanuel. *Fundamental Principles of the Metaphysic of Morals*. Translated by Thomas K. Abbott. Indianapolis: The Bobbs-Merrill Company, Inc., 1949.
- Kant, Immanuel. "On a Supposed Right to Lie from Altruistic Motives." 1-4. Accessed via PhilPapers June 6, 2017 at <http://www.mc.maricopa.edu/~davpy35701/text/kant-sup-right-to-lie.pdf>.
- Kirk, Robert. "Zombies." *The Stanford Encyclopedia of Philosophy*. Spring 2019. <https://plato.stanford.edu/entries/zombies/>.
- Kissinger-Knox, Alicia, Patrick Aragon, and Moti Mizrahi. "'Ought Implies Can,' Framing Effects, and 'Empirical Refutations.'" forthcoming in *Philosophia* 46, no. 1 (2018): 1-32.
- Knobe, Joshua, and Shaun Nichols. "An Experimental Philosophy Manifesto." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 3-14. New York: Oxford University Press, 2008.

- Knobe, Joshua. "Free Will and the Scientific Vision." In *Current Controversies in Experimental Philosophy*, edited by Elizabeth O'Neill and Eduard Machery. New York: Routledge, 2014.
- Knobe, Joshua, and Shaun Nichols, eds. *Experimental Philosophy*, New York: Oxford University Press, 2008.
- Kogut, Tehila and Ilana Ritov. "The Singularity Effect of Identified Victims in Separate and Joint Evaluations." *Organizational Behavior and Human Decision Processes* 97 (2005): 106-116. doi: 10.1016/j.obhdp.2005.02.003.
- Kurthy, Miklos, and Holly Lawford-Smith. "A Brief Note on the Ambiguity of 'Ought'. Reply to Moti Mizrahi's 'Ought, Can, and Presupposition: An Experimental Study.'" *Method* 4, no. 6 (2015): 244-249.
- Kurthy, Miklos, Holly Lawford-Smith, and Paulo Sousa. "Does Ought Imply Can?" *PLoS ONE* 12, no. 4 (2017): 1-24. <https://doi.org/10.1371/journal.pone.0175206>.
- Kvale, Steinar. *InterViews: An Introduction to Qualitative Research Interviewing*. Thousand Oaks: Sage Publications, 1996.
- Leben, Derek. "In Defense of 'Ought Implies Can.'" In *Oxford Studies in Experimental Philosophy, Volume 2*, edited by Tania Lombrozo, Joshua Knobe, and Shaun Nichols, chap. 6, Oxford: Oxford University Press, 2018. Kindle.
- "The Life You Can Save." The Life You Can Save. Accessed June 1, 2019. <https://www.thelifeyoucansave.org/>.
- May, Joshua. "Emotional Reactions to Human Reproductive Cloning." *Journal of Medical Ethics* 42 (2016): 26-30. doi: 10.1136/medethics-2015-102738.
- McGinn, Michelle K. "Transcript." In *The SAGE Encyclopedia of Qualitative Research Methods*, edited by Lisa M. Given, 1-4. Thousand Oaks: SAGE Publications, Inc., 2008. <http://dx.doi.org/10.4135/9781412963909.n462>.
- Mizrahi, Moti. "Ought, Can, and Presupposition: An Experimental Study." *Method* 4, no. 6 (2015): 232-243.
- Mizrahi, Moti. "Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith." *Method* 4, no. 6 (2015): 250-256.
- Nahmias, Eddy, Stephen G. Morris, Thomas Nadelhoffer, and Jason Turner. "Is Incompatibilism Intuitive?" In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 81-104. New York: Oxford University Press, 2008.

- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner. "Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility." *Philosophical Psychology* 18, no. 5 (2005): 561-584. doi: 10.1080/09515080500264180.
- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner. "The Phenomenology of Free Will." *Journal of Consciousness Studies* 11, no. 7-8 (2004): 162-179.
- Nichols, Shaun and Joshua Knobe. "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 105-126. New York: Oxford University Press, 2008.
- O'Connor, Timothy and Christopher Franklin. "Free Will." *The Stanford Encyclopedia of Philosophy*. Summer 2019. <https://plato.stanford.edu/entries/freewill/>.
- O'Neill, Elizabeth and Edouard Machery. "Introduction: Experimental Philosophy: What is it Good For?" In *Current Controversies in Experimental Philosophy*, edited by Elizabeth O'Neill and Edouard Machery. New York: Routledge, 2014.
- PhilPapers. "Experimental Philosophy." Edited by James Andow. Accessed May 30, 2019. <https://philpapers-org.ccl.idm.oclc.org/browse/experimental-philosophy>.
- PhilPapers. "Experimental Philosophy." Edited by Wesley Buckwalter. Accessed January 2, 2017. <https://philpapers.org/browse/experimental-philosophy>.
- PhilPapers. "Imaginative Resistance." Edited by Amy Kind. Accessed April 8, 2019. <https://philpapers.org/browse/imaginative-resistance>.
- PhilPapers. "The Gettier Problem." Edited by Ian M. Church. Accessed May 11, 2018. <https://philpapers.org/browse/the-gettier-problem>.
- Rothbauer, Paulette M. "Triangulation." In *The SAGE Encyclopedia of Qualitative Research Methods*, edited by Lisa M. Given, 1-4, Thousand Oaks: SAGE Publications, Inc., 2008. <http://dx.doi.org.ccl.idm.oclc.org/10.4135/9781412963909.n468>.
- Sarkissian, Hagop, Amita Chatterjee, Felipe De Brigard, Joshua Knobe, Shaun Nichols, and Smita Sirker. "Is Belief in Free Will a Cultural Universal?" *Mind & Language* 25, no. 3 (2010): 346-358. doi: 10.1111/j.1468-0017.2010.01393.x.
- Simone, Alina. *Experimental Philosophy Anthem*. Performed by Alina Simone. 2007; YouTube, <https://www.youtube.com/watch?v=tt5Kxv8eCTA>.
- Sinnott-Armstrong, Walter. "Abstract + Concrete = Paradox." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 209-230. New York: Oxford University Press, 2008.

- Sinnott-Armstrong, Walter. “‘Ought’ Conversationally Implies ‘Can.’” *The Philosophical Review* 93, no. 2 (1984): 249-261.
- Sinnott-Armstrong, Walter. “‘Ought to Have’ and ‘Could Have.’” *Analysis* 45, no. 1 (1985): 44-48.
- Sommers, Tamler. “Experimental Philosophy and Free Will.” *Philosophy Compass* 5, no. 2 (2010): 199-212. doi: 10.1111/j.1747-9991.2009.00273.x.
- Streumer, Bart. “I Ought to Reply, So I Can.” *Philosophia* (2018): 1-8. <https://doi.org/10.1007/s11406-018-0042-2>.
- Supplementary material for experiments 1, 2, and 3, from Vladimir Chituc et al., “Blame, Not Ability, Impacts Moral ‘Ought’ Judgments for Impossible Actions: Toward an Empirical Refutation of ‘Ought’ Implies ‘Can.’” *Cognition* 150 (2016): 20-25, accessed June 5, 2017, doi: 10.1016/j.cognition.2016.01.013. Supplementary material available at <http://dx.doi.org/10.1016/j.cognition.2016.01.013>.
- Sytsma, Justin. “The Robots of the Dawn of Experimental Philosophy of Mind.” In *Current Controversies in Experimental Philosophy*, edited by Elizabeth O’Neill and Eduard Machery. New York: Routledge, 2014.
- “The World Factbook: Denmark.” *Central Intelligence Agency: Library*. Last updated April 3, 2019. <https://www.cia.gov/library/publications/the-world-factbook/geos/da.html>.
- Thornton, Louise K., Amanda L. Baker, Martin P. Johnson, Frances Kay-Lambkin, and Terry J. Lewin. “Reasons for Substance Use Among People With Psychotic Disorders: Method Triangulation Approach.” *Psychology of Addictive Behaviors* 26, no. 2 (2012): 279-288. doi: 10.1037/a0026469.
- Tuggy, Dale. “Trinity.” *The Stanford Encyclopedia of Philosophy*. Winter 2016. <https://plato.stanford.edu/entries/trinity/>.
- Turri, John and Peter Blouw. “Excuse Validation: A Study in Rule-breaking.” *Philosophical Studies* 172, no. 3 (2015): 615-634. doi: 10.1007/s11098-014-0322-z.
- Van Someren, Maarten W., Yvonne F. Barnard, and Jacobijn A. C. Sandberg. *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes*. London: Academic Press, 1994.
- Von der Mühlen, Sarah, Tobias Richter, Sebastian Schmid, Elisabeth Marie Schmidt, and Kirsten Berthold. “Judging the Plausibility of Arguments in Scientific Texts: A Student–Scientist Comparison.” *Thinking & Reasoning* 22, no. 2 (2016): 221-249. Accessed June 7, 2017. doi: 10.1080/13546783.2015.1127289.

- Wagner-Tsukamoto, Sigmund. "Consumer Ethics in Japan: An Economic Reconstruction of Moral Agency of Japanese Firms – Qualitative Insights from Grocery/Retail Markets." *Journal of Business Ethics* 84, no. 1 (2009): 29-44. Accessed January 2, 2017. doi: 10.1007/s10551-008-9671-x.
- Wang, Yixuan. "The Mystery Revealed—Intersectionality in the Black Box: An Analysis of Female Migrants' Employment Opportunities in Urban China." *Hypatia: A Journal of Feminist Philosophy* 30, no. 4 (2015): 862-880. Accessed January 2, 2017. doi: 10.1111/hypa.12171.
- Weinberg, Jonathan M., Shaun Nichols, and Stephen P. Stich, "Normativity and Epistemic Intuitions." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 17-45. New York: Oxford University Press, 2008.
- Weiss, Zach. Lecturer. *A Survey of Experimental Philosophy and Epistemology*. YouTube, 15:21, April 26, 2015. <https://www.youtube.com/watch?v=YMDaoudmoEU>.
- Womack, Catherine and Norah Mulvaney-Day. "Feminist Bioethics Meets Experimental philosophy: Embracing the Qualitative and Experiential." *The International Journal of Feminist Approaches to Bioethics* 5, no. 1 (2012): 113-132. Accessed December 22, 2016. doi: 10.2979/intjfemappbio.5.1.113.
- Zuckerman, Phil. *Society Without God: What the Least Religious Nations Can Tell Us About Contentment*. New York: New York University Press, 2008.