

10-1-2000

Gender as an Impediment to Labor Market Success: Why Do Young Women Report Greater Harm?

Heather Antecol
Claremont McKenna College

Peter Kuhn

Recommended Citation

Antecol, Heather, and Peter Kuhn. "Gender as an Impediment to Labor Market Success: Why Do Young Women Report Greater Harm?" *Journal of Labor Economics* 18.4 (2000): 702-728. DOI: 10.1086/209974

This Article is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

NORC at the University of Chicago
The University of Chicago

Gender as an Impediment to Labor Market Success: Why Do Young Women Suffer?
Harm?

Author(s): Heather Antecol and Peter Kuhn

Reviewed work(s):

Source: *Journal of Labor Economics*, Vol. 18, No. 4 (October 2000), pp.

Published by: [The University of Chicago Press](#) on behalf of the [Society of Labor Economists](#)

[NORC at the University of Chicago](#)

Stable URL: <http://www.jstor.org/stable/10.1086/209974>

Accessed: 09/11/2012 16:08

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and access the full text of articles, books, and other content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press, Society of Labor Economists, NORC at the University of Chicago are collaborating with JSTOR to digitize, preserve and extend access to *Journal of Labor Economics*.

<http://www.jstor.org>

Gender as an Impediment to Labor Market Success: Why Do Young Women Report Greater Harm?

Heather Antecol, *Illinois State University*

Peter Kuhn, *University of California, Santa Barbara*

Compared to older women, young female job seekers are more than three times as likely to report that their ability to find a good new job is compromised by their gender. This phenomenon cannot be statistically attributed to observed personal or job characteristics, or to any “objective” measure of discrimination. Further, women’s reports of gender-induced *advantage*, and men’s reports of gender-induced harm, are also more prevalent among the young. A possible interpretation of all these patterns is that young people are more likely to interpret a given departure from gender-neutral treatment as causally affected by their gender.

I. Introduction

In a number of recent articles (e.g., Kuhn 1987, 1990; Barbezat and Hughes 1990; Heywood 1992; Hampton and Heywood 1993, 1996; Laband and Lentz 1993, 1998; Johnson and Neumark 1997), labor economists have begun to analyze aspects of labor market discrimination that, unlike more standard “residual wage gap” measures (Oaxaca 1973), are directly derived from survey reports of discriminatory treatment. Survey

The authors would like to thank Human Resources Development Canada for making the data for this study available and the Canadian International Labour Network (CILN) for financial support. CILN receives major funding from the Social Sciences and Humanities Research Council of Canada and McMaster University. Felice Martinello and David Robinson provided helpful comments.

[*Journal of Labor Economics*, 2000, vol. 18, no. 4]
© 2000 by The University of Chicago. All rights reserved.
0734-306X/2000/1804-0006\$02.50

reports of discrimination can be of interest to economists for a number of reasons, including as a check on the validity of traditional “wage gap” measures, and as objects of interest in their own right. For a number of issues, including public support for antidiscrimination policies, and the prevalence of the kind of discriminatory treatment that may be remediable in courts of law, survey reports of discrimination may in fact be more relevant than wage gap measures.¹

A common, but often overlooked, finding in the literature on survey reports of discrimination is a greater tendency for younger women to feel harmed by discrimination than older women (see, e.g., Kuhn 1987; and Laband and Lentz 1998). What explains this? Because measured gender-wage gaps tend to be lower among young women than among older women, earlier work has tended to attribute higher reports of discrimination among young women to a larger amount of unmeasured discriminatory treatment.² In this article, we shed new light on this and other hypotheses about reported discrimination using a new survey of Canadian job seekers. This survey contains more extensive information on perceived discrimination than has been available in earlier surveys; in particular, it contains questions on gender-induced advantage as well as harm and asks the reported gender discrimination questions of both men and women.³

We begin our analysis by replicating earlier results: in our data, female job seekers aged 15–24 are more than three times as likely to say they have been hurt by their gender in the labor market than those aged 55–64. We then assess various possible explanations of this phenomenon in turn. We show, first, that higher reports of harm among young women cannot be explained by any measured personal or job characteristics. Second, these reports cannot be explained by any standard measures of “objective” discrimination that have been used by economists: in our sample, measured discrimination on a very wide array of outcome measures is always lower among young women.

Third, using the extra information in the current survey on women’s perceptions of gender-induced advantage, and on men’s perceptions of

¹ While econometric evidence of wage gaps is becoming increasingly accepted in many courts, evidence of specific incidents of discriminatory treatment—which may or may not be reflected in wage measures—is still seen as stronger support for a claim of discrimination in almost all cases.

² For example, Kuhn (1987, p. 581) concludes that “young, well-educated women behave as if they experience considerably more nonstatistical evidence of discrimination than others.”

³ The nature of our survey allows us as well to examine the effects on reports of a much wider array of statistical measures of discrimination: in addition to wage gaps, we are able to consider gaps in unemployment, wage changes between jobs, and reservation wages.

gender-induced harm, we are able to rule out a higher mean level of unmeasured discrimination against young women as an explanation of their more frequent reports of harm. The reason for this is a robust and surprising “symmetry” that we find in the age patterns of reported harm and advantage: for both women and men, the frequency of both reported harm and advantage declines with age. Clearly, these patterns cannot be explained by a higher overall level of discrimination against young women, even of a form that is not captured by standard discrimination measures.

Finally, we estimate a formal model of the decision to report gender-induced advantage or harm that allows us to test two further explanations of the above patterns. One of these—greater dispersion in the labor market experiences of young workers—receives mixed support: essentially, if our result is driven by age-related differences in dispersion, these differences must be of a form that is not captured by any of the measured characteristics in our sample, including education, marital status, presence of young children, job tenure, and percent female in one’s occupation. The other explanation involves a particular kind of age difference in reporting behavior. According to this interpretation, young workers of both sexes are simply more willing to label departures in either direction from gender-neutral treatment as being caused by their gender. In this sense, young workers are less “tolerant” of both traditional and reverse discrimination, a phenomenon that might pose a challenge for the design of future antidiscrimination policy.

Section II describes the data used in this study. Section III presents our main results concerning women’s reports of gender-induced harm. Section IV broadens the analysis to include women’s reports of gender-induced advantage, and men’s reports of harm, and estimates our formal model of reporting. Section V concludes.

II. Data

The data used in our analysis are drawn from a nationally representative sample of Canadians who have recently experienced a job separation: the 1995 Canadian Out of Employment Panel (COEP). Individuals separating from jobs in 2 window periods during 1995 were identified using administrative records of the Unemployment Insurance system, which requires employers to file a Record of Employment (ROE) form whenever a separation occurs. The data contain a rich set of measures about a worker’s preseparation job, his or her first postseparation job, the (post-separation) job at the time of the interview, and on unemployment spells; our analysis focuses on the actual and perceived outcomes of these workers’ search for a new job. Unlike some other studies of discrimination, the current study thus focuses particularly on actual and perceived discrimination in access to jobs, rather than, for example, discrimination

within jobs. This is useful, in our view, because it is widely accepted that the former issue of access to jobs, rather than the latter, plays the central role in explaining male-female wage gaps today (see, e.g., Johnson and Solon 1986).⁴

Because of a problem with how reports of gender-induced harm were measured, in this article we only use the information from separations in the first window of the 1995 COEP survey, which consists of 3,898 individuals.⁵ Eliminating individuals who were 65 years of age or over left us with a sample of 1,586 women and 2,280 men. Descriptive statistics on the main variables used in our analysis are presented in table 1. Inspection of table 1 reveals that, on average, women in our sample are slightly older than men but have about 10 weeks less tenure on their preseparation jobs. Further, women are more likely to have education at the higher levels (college and university) compared to the men in our sample. Not surprisingly, women are much more likely to be in highly feminized occupations, though they do not differ markedly from men in their marital status distribution and presence of children. Turning to labor market “outcome” variables, table 1 shows that the mean hourly preseparation wages of men and women are \$15.72 and \$10.09, respectively. Thus, women earn 64% of what men do before the separation. A slightly smaller earnings ratio (61%) is found for postseparation wages. In our sample of unemployment durations, women had been unemployed about 3 weeks longer than men as of the survey date, which was usually about 38 weeks after the separation.⁶

Finally, the measure of gender-induced harm in the current article is based on the following question: “In any of the job search that you have done since [the separation date], do you feel that your gender has had an impact on your ability to find a good job?” To avoid framing the question in a way that might encourage responses in either direction, the allowed

⁴ The focus on job seekers also serves to eliminate a potential explanation of the age difference in reports suggested in Kuhn (1987): the notion that young women might have had less of a chance to locate a nondiscriminatory employer than older women. In the current data, women of all ages are engaged in the search for a new job, yet we continue to see the patterns found in nationally representative surveys of reported discrimination.

⁵ In cohort 2 the question on perceived harm was only asked of people who, at the survey date, were still searching for a job. To the extent that these individuals are still searching because they have had disappointing search outcomes, or because they can afford to search longer than others, they will be systematically different from the population of all job separations.

⁶ In some of what follows we analyze the unemployment durations of the individuals in our sample. When we do so our sample is further restricted to individuals who lost their job due to a layoff, and who reported that they engaged in at least some search for a new job after the layoff. See n. 14 below.

Table 1
Sample Mean Characteristics

	Men		Women	
	Mean	(SE)	Mean	(SE)
Age	36.65	(.23)	37.41	(.27)
Education variables:				
Less than high school	.34	(.01)	.25	(.01)
High school	.31	(.01)	.32	(.01)
Some college or university	.14	(.01)	.13	(.01)
College	.12	(.01)	.16	(.01)
University	.10	(.01)	.14	(.01)
Marital status:				
Married	.61	(.01)	.62	(.01)
Separated/divorced/widowed	.08	(.01)	.12	(.01)
Single	.31	(.01)	.26	(.01)
Children less than 6 years old	.28	(.01)	.23	(.01)
Percent female in occupation	.25	(.00)	.55	(.01)
Tenure (weeks)	230.53	(7.79)	219.79	(7.19)
Preseparation hourly wage	15.72	(.16)	10.09	(.13)
Postseparation hourly wage	16.37	(.19)	10.06	(.16)
Unemployment spell (weeks)	20.70	(.44)	23.84	(.68)
Discrimination measures:				
Hurt	.06	(.00)	.10	(.01)
Help	.05	(.00)	.04	(.00)
Number of observations	2,280		1,586	

NOTE.—Unemployment spells include incomplete spells. Due to missing data, the number of observations is lower for some variables.

responses were (1) yes, hurt; (2) yes, helped; or (3) no impact. In cohort one, which forms the basis of our sample, the question was asked of all individuals, irrespective of gender, and irrespective of their employment status at the time of the interview. For brevity in what follows, we occasionally refer to this as the “reported discrimination” question; to individuals choosing the “yes, hurt” response as experiencing “hurtful discrimination”; and to those indicating “yes, helped” as experiencing “helpful discrimination,” that is, being a beneficiary of labor market discrimination.⁷

According to table 1, about 14% of women and 11% of men experiencing a job separation report that their gender had some effect on their ability to find a good new job, with the balance—a vast majority of both men and women—indicating they felt their gender had no effect. Among those who reported advantage or harm, women were more likely to feel that they were hurt than helped, by a ratio of about 10 to 4, while men’s

⁷ It may be objected that this question does not include the word discrimination. Although this is true, the age patterns of reported “discrimination” in our study are similar to a number of other studies (e.g., Laband and Lentz 1998), increasing our confidence that they capture the same phenomenon.

reports were almost evenly split between those who were hurt or helped. Thus, while hurtful discrimination against women is the most common way in which gender is perceived to affect search outcomes, other forms of perceived gender-induced harm or advantage also play substantial roles in our sample.

III. Women's Reports of Gender-Induced Harm

Table 2 shows the fraction of women reporting gender-induced harm in five age categories. As in Kuhn (1987), it is clear that women's reports of hurtful discrimination are highest among the young: in the present case they are more than three times as frequent among women under 25, at 13.2% of the sample, than among women aged 55–64, at 4.3%. This difference is the main stylized “fact” we analyze in this article and echoes earlier findings in Kuhn (1987) and others.

Can young women's more frequent reports of gender-induced harm be explained by differences in observed characteristics between them and older women? To answer this question, we first estimate a probit model of hurtful discrimination on age only. We then add observed covariates to see if, under any specification, the direct effect of age is substantially reduced in magnitude. Before discussing our results, it is worth commenting on the role of one particular variable included in these regressions: the gender composition of one's occupation.⁸ We included this variable to capture what we initially considered a very appealing explanation for the pattern observed in our data: perhaps young women are more likely to be “pioneers,” that is, be among the first women to enter high-paying, formerly all-male occupations such as law, science, medicine, and management. High levels of reported discrimination could occur because these workplaces had not yet adapted themselves to women's presence and because of lingering sexist attitudes; this could be consistent, as noted below, with low levels of measured discrimination because these are high-paying occupations for women.⁹

⁸ Percent female was calculated for the 19 occupations coded in the COEP survey from a nationally representative sample of workers in the 1994 Survey of Consumer Finances (SCF). Like all the right-hand-side variables used here, this refers to the pre-separation job to avoid endogeneity. The correct interpretation is, we believe, that women with previous jobs in highly feminized occupations are likely to have skills that are specific to such occupations. Their search activity is therefore more likely to focus on those kinds of jobs.

⁹ A formal version of the “pioneers” hypothesis can be found in Kuhn (1993). In that model, binding entry restrictions against women seeking access to jobs “designed” for men become more likely as women's labor force attachment becomes more similar to men's, and as the gender wage gap consequently narrows. (The intuition is that women will not want those jobs until they become sufficiently committed to the work force.) Thus, to the extent that wanting a

Table 2
Women's Reports of Gender-Induced
Harm, by Age Group

Age	Gender-Induced Harm (Fraction Reporting)
55-64	.043 (.005)
45-54	.088 (.007)
35-44	.093 (.007)
25-34	.123 (.008)
15-24	.132 (.008)

NOTE.—Exact binomial standard errors in parentheses.

Table 3 presents estimates of various specifications of a probit for reported gender-induced harm. The following results are of note. Although the age coefficient does tend to decline in significance as controls are added, it remains negative and of roughly the same magnitude no matter what the specification. In particular, the difference in the predicted probability of reporting gender-induced harm between women aged 20 and women aged 50, holding all other characteristics fixed at their mean level, ranges from 0.144 to 0.097. The only other variables that affect reports of hurtful discrimination are the percent female in one's occupation and, perhaps, marital status. Percent female does have the effect anticipated on reports: women who are a minority in their occupations (low values of percent female) are much more likely to report encountering hurtful discrimination in their search for a new job, perhaps simply because they are more likely to be interacting with, and competing with, men in the job-search process. Importantly, however, this does not help explain the age pattern in reports in our data: adding the percent female variable has essentially no effect on the estimated age effect on reported harm. The reason is that, in our data, age and percent female are hardly related: while "pioneers" do run into more barriers, apparently women of all age groups have their share of pioneers.

We conclude from table 3 that, while some of the observable characteristics of women in our sample have interesting and potentially important effects on women's reports of gender-induced harm, none of these other characteristics is able to account for the observed effect of age on reported harm. Therefore, we examine an alternative explanation for

traditionally "male" job but being denied entry to it is a kind of discrimination, this kind of discrimination will (a) not be captured by standard wage-gap measures and (b) be most common among those groups of women (e.g., younger cohorts) who are most committed to the labor market and face the smallest measured wage gaps.

Table 3
Probit Coefficients for Gender-Induced Harm, Various Specifications

	(1)	(2)	(3)	(4)	(5)
Age	-.0134 (3.20)	-.0138 (3.15)	-.0115 (2.05)	-.0112 (1.99)	-.0106 (1.85)
High school		-.0142 (.12)	.0098 (.08)	-.0171 (.14)	.0151 (.12)
Some college/university		-.1490 (.95)	-.1758 (1.09)	-.1807 (1.10)	-.1173 (.71)
College degree		-.1387 (.93)	-.1442 (.95)	-.1251 (.81)	-.0445 (.28)
University degree		.1932 (1.38)	.1953 (1.36)	.1747 (1.19)	.2472 (1.64)
Married			-.2465 (2.17)	-.2186 (1.89)	-.2270 (1.96)
Separated, divorced, widowed			.1366 (.85)	.1494 (.92)	.1224 (.75)
Children less than 6 years old			-.0396 (.45)	-.0472 (.52)	-.0409 (.45)
Tenure			-.0001 (.37)	.0000 (.20)	.0000 (.05)
Percent female in occupation					-.5660 (2.72)
Province dummies	no	no	no	yes	yes
Number of observations	1,548	1,519	1,473	1,471	1,466

NOTE.—Absolute value of z-statistics in parentheses.

young women's more frequent reports of hurtful discrimination: higher levels of measured labor market discrimination against them. To examine this hypothesis, we first compute individual-specific measures of discrimination for three labor market outcomes—postseparation wages, preseparation wages, and unemployment durations, as follows.¹⁰ We compute three alternative measures because, especially in this sample of job searchers, there may be a number of plausible ways in which women are affected by discrimination, and we want to be sure that no plausible channel by

¹⁰ Most plausibly, we would probably expect reported discrimination during the job search interval to be related to the quality of the job that was ultimately found (i.e., the postseparation wage) and the length of time required to find it (unemployment duration). However it is also possible that the frequency of discriminatory treatment encountered by a population subgroup (e.g., young women) can have equilibrium effects on the wages or unemployment durations of all members of the group (e.g., via a lower reservation wage in the search for the preseparation as well as the postseparation job). If it does, then the preseparation wage—which has the advantage of being observed for substantially more individuals in our data—should also give an indication of the distribution of discrimination across broad population subgroups.

which measured discrimination patterns might explain the age pattern in reports is overlooked.¹¹

For the case of wage discrimination, we compute our measure of “objective” discrimination by first estimating the following log wage regressions:

$$w_i^F = \sum_{j=0}^J X_{ij}^F \beta_j^F + \mu_i^F \quad (1)$$

and

$$w_i^M = \sum_{j=0}^J X_{ij}^M \beta_j^M + \mu_i^M, \quad (2)$$

where F stands for female, M stands for male, w_i^F and w_i^M are individual log wages, and X_{ij} is a vector of ones for $j = 0$. We then use the estimated coefficients from equations (1) and (2) to compute two alternative individual-specific measures of discrimination against women, as follows:¹²

$$\hat{D}_i^1 = \sum_{j=0}^J X_{ij}^F \hat{\beta}_j^M - w_i^F \quad (3)$$

and

$$\hat{D}_i^2 = \sum_{j=0}^J X_{ij}^F \hat{\beta}_j^M - \sum_{j=0}^J X_{ij}^F \hat{\beta}_j^F. \quad (4)$$

Conceptually, both the above definitions of discrimination can be thought of as the (log) difference between what a specific woman actually

¹¹ In our analysis we also looked at two other outcomes: wage differences between the preseparation and postseparation jobs and reported reservation wages. Consistent with our prior expectation that the amount of discrimination is a relatively permanent characteristic of the market for workers of a given age, wage differences were unrelated to age, while reservation wages followed the same pattern as preseparation and postseparation wage levels. Thus neither of these are able to explain the age pattern in reports either.

¹² More commonly, discrimination measures have been computed as a predicted log wage difference for an “average” woman in the sample (Oaxaca 1973). The measures used here simply apply Oaxaca’s decomposition to each individual woman in the sample (see Kuhn 1987).

earns and what she would earn “if she were a man.” The difference between the two concerns the definition of the counterfactual: \hat{D}_i^2 implicitly compares a woman to a man with both her observed characteristics and unobserved ability (thus the earnings function residual drops out of the expression); \hat{D}_i^1 compares a woman’s actual earnings to a man with her observed characteristics of “average” unobserved ability.¹³ Because, even controlling for measurable characteristics, men tend to earn more than women, we expect the majority of individual \hat{D}_i ’s, according to either measure, in our sample to be positive.

Our measure of discrimination in unemployment durations is computed as similarly to the wage measures as possible, given the nature of our data. We first estimate male and female log duration regressions in a manner strictly parallel to (1) and (2), using a censored-normal model to account for incomplete spells. Unlike the wage regressions, however, this regression was restricted to individuals who lost their job due to a layoff and who engaged in at least some job search after the separation.¹⁴ Because longer durations are worse labor market outcomes than shorter ones, our estimates of discrimination are then the negatives of (3) and (4). The measure, \hat{D}_i^1 , that relies on actual individual durations is treated as missing for individuals whose durations are censored.

In table 4, we present unadjusted means of the three different labor market outcomes, and of individual-specific discrimination measures based on each outcome, for women in different age categories. In all the regressions underlying table 4, the following variables are included in the X_{ij} ’s: age, education, region, marital status, number of children less than 6 years of age, and tenure on the preseparation job.¹⁵ Two immediately apparent features of table 4 are the similarity of the patterns in unadjusted

¹³ A priori, it is not clear which of these two comparisons we would expect most women to make. To the extent that individuals observe more about each other than the econometrician, the former comparison (a man with similar observed and unobserved ability) would seem more appropriate. On the other hand, individuals might discount components of either good or bad “luck” in their own earnings residuals when making these comparisons, which tends to favor the latter measure, i.e., \hat{D}_i^1 .

¹⁴ We determined who was actively involved in search based on two questions in the COEP survey: “Did you look for work between the separation date and the first job [you held since the separation]?” (only asked of people who had a first job) and “Did you look for work between the separation date and [the survey date]?” (only asked of people who had no jobs since the separation). If the answer was no to either of these questions, the respondents were dropped from the sample.

¹⁵ Percent female is not included in the calculation of measured discrimination because we adopt the view that women’s concentration in highly feminized occupations is an outcome of discrimination. That said, the results do not differ markedly when it is included.

Table 4
Selected Measures of “Objective” Discrimination, by Age Group

Age	Measured Discrimination: Various Outcome Measures								
	Preseparation Wages			Postseparation Wages			Unemployment Durations*		
	Unadjusted Gap (1)	Adjusted Gap (\hat{D}^1) (2)	Adjusted Gap (\hat{D}^2) (3)	Unadjusted Gap (4)	Adjusted Gap (\hat{D}^1) (5)	Adjusted Gap (\hat{D}^2) (6)	Unadjusted Gap (7)	Adjusted Gap (\hat{D}^1) (8)	Adjusted Gap (\hat{D}^2) (9)
55-64	.569	.738	.602	.660	.831	.678	.358	.976	.377
45-54	.575	.564	.556	.609	.633	.615	.392	.551	.292
35-44	.504	.438	.494	.546	.493	.553	.468	.608	.317
25-34	.372	.340	.377	.423	.403	.435	.172	.476	.280
15-24	.259	.450	.322	.373	.504	.405	.071	.791	.031

* Female unemployment duration minus male. Censoring adjusted for using a censored normal regression model of log spell duration, separately by age category.

versus regression-adjusted wage gaps, and the similarity of patterns for preseparation and postseparation wages. In one sense, neither of these should be surprising: measured discrimination is a relatively permanent characteristic of the labor market for workers of a given age and education level, and (because men's and women's characteristics differ little) it is largely driven by gender differences in the unadjusted wage gap. Further, unlike survey reports of discrimination, all three estimates of "objective" discrimination against women are lower among young workers, though the patterns are stronger and more consistent for wages than unemployment spells. (There is also an exception for the youngest age category when wage discrimination is calculated using \hat{D}_i^1 .) Although we explore these issues further in the multivariate analysis below, we draw two main conclusions from these patterns. First, given these raw correlations, patterns of measured discrimination are not likely to be able to explain patterns of reported gender-induced harm across age groups. Second, the much weaker age pattern in unemployment than wage gaps suggests the following interpretation of our evidence: compared to men of their own age, young women face a more favorable wage distribution than older women and take advantage of this by raising their reservation wages in searching for both their preseparation and postseparation jobs. This is further justification for our treatment in this article of both postseparation and preseparation wage levels as outcomes of the general level of demand, supply, and discrimination conditions in labor markets for workers in a given age group.

To get a more formal idea of the ability of measured discrimination to explain the age pattern of reported discrimination, we estimate probit models analogous to those presented in table 3 adding controls for "observed" discrimination. Panels A and B of table 5 present the estimates of various specifications of the probits of hurtful discrimination when discrimination is calculated for postseparation wages using the measures \hat{D}_i^1 and \hat{D}_i^2 , respectively. The following results are of note. As in table 3, the age coefficient again fluctuates in significance but remains negative and of roughly the same magnitude no matter what the specification. In particular, the difference in the predicted probability of reporting gender-induced harm between women aged 20 and women aged 50, holding all other characteristics fixed at their mean level, ranges from 0.177 to 0.112 in panel A and ranges from 0.248 to 0.097 in panel B. Additionally, percent female in one's occupation continues to be a negative and significant determinant of reports. Finally, it is worth noting the absence of a robust relationship between measured discrimination and women's survey reports of discrimination. The \hat{D} coefficients in panels A and B in table 5 are often negative, and in all cases but one are highly insignificant, further supporting the notion that measured discrimination is of little

Table 5
Probit Coefficients for Gender-Induced Harm, Controlling for “Objective”
Discrimination on the Postseparation Job

	(1)	(2)	(3)	(4)	(5)
A. Controlling for \hat{D}^1 :					
Age	-.0182 (3.12)	-.0202 (3.33)	-.0155 (2.09)	-.0144 (1.92)	-.0146 (1.92)
\hat{D}^1	-.0031 (.02)	.0100 (.08)	.0397 (.32)	-.0024 (.02)	-.0086 (.07)
Percent female in occupation					-.5808 (2.11)
Education	no	yes	yes	yes	yes
Demographics	no	no	yes	yes	yes
Province	no	no	no	yes	yes
Number of observations	927	927	927	908	907
B. Controlling for \hat{D}^2 :					
Age	-.0116 (2.37)	-.0157 (2.78)	-.0216 (3.00)	-.0112 (1.99)	-.0106 (1.85)
\hat{D}^2	-.1326 (.48)	.3342 (.65)	1.507 (2.23)	dropped*	dropped*
Percent female in occupation					-.5660 (2.72)
Education	no	yes	yes	yes	yes
Demographics	no	no	yes	yes	yes
Province	no	no	no	yes	yes
Number of observations	1,471	1,471	1,471	1,471	1,466

NOTE.—In these tables \hat{D}^2 was calculated for all workers—even those without postseparation jobs—from the postseparation log wage regressions. Very similar results are obtained if the sample is restricted to those with postseparation jobs. Absolute value of z-statistics in parentheses.

* We exclude \hat{D}_i^2 from the estimating equation because \hat{D}_i^2 is a linear combination of the other included variables.

help in explaining patterns of reported discrimination across age groups, or across individuals.

Table 6 assesses the robustness of the results in panels A and B in table 5 by presenting age coefficients from probits analogous to those in table 5, but when “objective” discrimination is measured, alternatively, using preseparation wages and unemployment durations.¹⁶ As can be seen, the age coefficients found in table 6 display the same general patterns as those found in panels A and B in table 5, although the patterns are weaker for unemployment spells. Further, percent female in occupation once again remains a negative and significant influence on reports. (There is an exception when unemployment duration discrimination is calculated using \hat{D}_i^2 .) Thus, the results found in panels A and B in table 5 seem to be robust.¹⁷

¹⁶ Similar results are also found when discrimination in reported reservation wages, wage changes, or employment rates are used in these regressions.

¹⁷ The one exception noted above likely results from collinearity between

Table 6
Probit Coefficients of Gender-Induced Harm on Age, Controlling for
Discrimination in Other Dimensions

Observed Characteristics Controlled for:	Preseparation Wage		Unemployment Spell	
	\hat{D}^1 (1)	\hat{D}^2 (2)	\hat{D}^1 (3)	\hat{D}^2 (4)
\hat{D} only	-.0123 (2.65)	-.0104 (2.13)	-.0317 (2.60)	-.0136 (1.70)
Education	-.0135 (2.83)	-.0118 (2.26)	-.0366 (2.85)	-.0138 (1.69)
Education, demographics	-.0122 (2.05)	-.0182 (2.44)	-.0287 (1.86)	-.0113 (1.05)
Education, demographics, province	-.0116 (1.93)	-.0112 (1.99)	-.0373 (2.25)	-.0145 (1.34)
Education, demographics, province, percent female in occupation	-.0107 (1.77)	-.0106 (1.85)	-.0421 (2.38)	-.0143 (1.31)
“Femocc” coefficient in above	-.7286 (3.33)	-.5660 (2.72)	-1.453 (2.24)	-.5863 (1.48)

NOTE.—Absolute value of z-statistics in parentheses.

One final, potential concern with our result in this section is “omitted experience bias.” In particular, because our measures of labor market experience are limited to age, education, and tenure on the preseparation job, unobserved differences in labor market experience between men and women are likely to exist. More importantly, they are likely to be greater among older than younger individuals, suggesting that our estimates of measured discrimination may be biased upward for older women.

While “omitted experience bias” may be an important issue in the precise quantification of gender-wage-gap measures, it can explain the higher levels of reported gender-induced harm among young women in our sample only if it is severe enough to reverse the observed positive correlation between measured discrimination and age. We feel this is highly unlikely for the following reasons: first, because labor market experience cannot differ greatly among young workers, measured discrimination among the youngest cohort in our sample is unlikely to be affected by omitted experience bias. This is the smallest amount of discrimination for any age group in our sample (see table 4), and “true”

personal characteristics and measured discrimination, which becomes substantial when the list of covariates comes close to exhausting the list of X variables used to calculate the \hat{D} measures.

discrimination faced by older women would have to be lower than this measure if omitted experience bias explains our results. Second, when we include a measure of “partial” labor market experience as an additional underlying determinant of measured discrimination, the magnitude and significance levels of the age coefficients presented in tables 5 and 6 are virtually unchanged.¹⁸

Third, unlike our data, the 1985 Panel Study of Income Dynamics (PSID) does collect information on total labor market experience. When we analyzed age patterns of measured wage discrimination in this data set, we found that controlling for actual labor market experience does not reverse the positive correlation between measured discrimination and age. In particular, the PSID asked the following questions of household “heads” and their “wives”: “How many years altogether have you worked for money since you were 18?” and “How many of these years did you work full-time for most or all of the year?” Using this information we are able to construct a measure of full-time experience, part-time experience, and time out of the labor force.¹⁹ We then reestimate the log wage regressions given by equations (1) and (2), where the following variables are included in the X_{ij} 's: years of education, full-time experience, full-time experience squared, part-time experience, part-time experience squared, years out of the labor force, marital status, number of children, racial dummy variables, and state dummy variables.²⁰ As before, we use the estimated coefficients from equations (1) and (2) to compute two alternative individual-specific measures of discrimination against women given by equation (3) and (4).

In appendix table A1, we present unadjusted means of log wages, and of individual-specific discrimination measures based on log wages, for women in different age categories using the 1985 PSID data. As in table

¹⁸ The COEP does contain information on whether a respondent had income from wages in any of the 5 years preceding the survey. Using this information, we constructed a “partial” measure of actual labor market experience equal to the number of years in the past 5 in which they had income from wages. All “partial” labor market experience results are available from the authors upon request.

¹⁹ Full-time experience is equal to the number of years the respondent reported working full-time since the age 18. Part-time experience is equal to the total number of years the respondent reported working since the age 18 minus the number of years of full-time experience. Finally, time out of the labor force is equal to the number of years the individual was out of the labor force since they graduated from school.

²⁰ Interestingly, we find that time out of the labor force has a negative and significant effect on the wages of women, but not of men. Further, men receive a higher return to each additional year of both full-time and part-time experience than women. In fact, women do not receive a wage premium for part-time work experience. All the remaining coefficients have the expected sign and are available from the authors upon request.

4, there is a similarity of the patterns in the unadjusted versus the regression-adjusted wage gaps. Further, measured discrimination is lower among younger workers than among older workers despite the inclusion of actual labor market experience as an additional determinant of measured discrimination.

The final and most important reason why omitted experience bias cannot explain our results is the evidence we present in the next section, that like gender-induced harm, survey reports of gender-induced advantage are also more common among young women. This is very hard to explain with a story based only on higher unmeasured discrimination faced by young women; clearly a more complex explanation than omitted experience alone is required.

We conclude the following from the analysis in this section. First, observable characteristics of the women in our sample are unable to account for the observed effect of age on reported discrimination. Second, neither can any standard measure of “objective” discrimination explain the more frequent reports of hurtful discrimination made by younger women in our sample.

IV. Broadening the Analysis: Gender-Induced Advantage and Men’s Perceptions

In this section we broaden our analysis by examining women’s reports of gender-induced advantage in the labor market, as well as men’s reports of gender-induced advantage and harm. The main goal is to ask whether this additional information can shed any more light on young women’s much greater propensity to report gender-induced harm. We begin our analysis in table 7, which presents means of these additional indicators by age. These means reveal what were to us some very surprising patterns: first, like women, young men are also more likely to feel that they have been hurt by gender discrimination than older men, though the relationship is certainly not as strong as it is for women.²¹ Further, young women, while more likely to report gender-induced harm, are also more likely than other women to report that they benefited from being female.²² Thus, especially for women, reports of both harm and advantage

²¹ Interestingly, men’s reporting patterns across education groups are however much stronger than women’s, but of a similar nature to women’s reporting patterns across age groups: highly educated men are much more likely to say they were hurt by discrimination, and more likely to say they were helped by discrimination, than less-educated men. We view this as a useful topic for further research.

²² We estimated a probit model of gender-induced advantage on age only for women. We then add observed covariates to see if, under any specification, the direct effect of age is substantially reduced in magnitude. Similar to our results on

Table 7
Frequency of Other Aspects of Reported Discrimination, by Age Group

Age	Women		Men		
	Help (1)	Hurt + Help (2)	Hurt (3)	Help (4)	Hurt + Help (5)
55-64	.011	.053	.045	.038	.083
45-54	.034	.122	.042	.030	.071
35-44	.027	.120	.058	.063	.122
25-34	.051	.174	.061	.050	.111
<25	.064	.195	.067	.058	.125

seem to move in tandem across age categories. This is illustrated in columns 2 and 5 of table 7, which simply add together all those individuals who reported either harm or advantage. In all cases, these fractions fall with age, generally more strongly than reports of harm or advantage alone. Estimating a probit model of reports of either harm or advantage on age (plus all additional control variables found in col. 5 of table 3), separately for men and women, indicates that younger individuals are more likely to depart from gender neutrality than older individuals, however, this relationship is only significant for women (See appendix table B1).²³

In order to interpret these patterns, in this section we develop a framework of the decision to report both gender-induced harm and advantage in a confidential survey. An interesting challenge for this framework arises from the fact that, while our dependent variable is clearly ordered (one's gender either hurt, had no impact, or helped), a worker's age appears (at least in our raw data) to have the effect of increasing the frequency of both "extreme" outcomes ("hurt" and "helped"). As most existing techniques for dealing with ordered data cannot fit such patterns well, we develop some simple models that are better suited to this purpose.

gender-induced harm, we find that, although the age coefficient does tend to decline in significance as controls are added, it remains negative and of roughly the same magnitude no matter what the specification.

²³ Although the age coefficient in the probit for reports of either harm or advantage on age is negative and significant for males, once controls for education are added, the significant, negative relationship ceases to exist. This is not surprising given n. 21 above, i.e., men's reporting patterns across education groups are much stronger than women's.

A. Conceptual Framework

Assume that the net amount of favorable treatment, relative to the opposite sex, faced by individual i in his or her job search is given by a scalar, Δ_i ,

$$\Delta_i = \sum_{j=0}^J \theta_j Z_{ij} + \phi_i \quad (5)$$

where ϕ_i is a normal error term with mean zero and variance σ_ϕ^2 . What we have in mind as entering into Δ_i includes longer or shorter unemployment spells relative to the opposite sex, differential treatment in job interviews, outright sexual harassment, and any other labor market outcome differentials encountered by the individual (whether observed or unobserved by the econometrician) relative to a similar person of the opposite sex. To allow for reports of both gender-induced harm and advantage, suppose that a woman reports that her job search was unaffected by her gender if and only if

$$-K < \Delta_i < K, \quad (6)$$

where K is the threshold amount of differential treatment that induces reports of either harm (in the lower tail of the distribution of Δ_i) or advantage (in the upper tail).²⁴

Combining (5) and (6), an individual reports neither harm or advantage if and only if

$$\left(-K - \sum_{j=0}^J \theta_j Z_{ij} \right) / \sigma_\phi < \mu_i < \left(K - \sum_{j=0}^J \theta_j Z_{ij} \right) / \sigma_\phi \quad (7)$$

or

$$\left(K_1 - \sum_{j=1}^J \theta_j Z_{ij} \right) / \sigma_\phi < \mu_i < \left(K_2 - \sum_{j=1}^J \theta_j Z_{ij} \right) / \sigma_\phi \quad (8)$$

²⁴ For now, we assume a constant reporting threshold across individuals; in what follows we shall explicitly model the dependence of this threshold on age. Note also that, because Δ_i has an intercept (see eq. [8]), (6) does not constrain reports of harm and advantage to be equally frequent; instead the symmetry of (6) is meant to capture the notion of individuals implicitly conducting two-tailed hypothesis tests about whether discrimination exists, which by construction are symmetric.

where μ_i is a standard normal variate. In (8), $K_1 = \theta_0 - K$, and $K_2 = \theta_0 + K$. Reports of hurtful and helpful discrimination are given by the obvious complementary expressions.

The model in (8), as written, is a standard ordered probit equation with three ordered responses. As is well known, σ_φ which in our model represents the standard deviation of individuals' experiences that is not associated with observed characteristics, is not identified; the standard practice is to normalize it to one and interpret the coefficient estimates as relative to this term, that is, as estimates of θ/σ_φ . Note that this model also generates estimates of standardized "cutoffs," K_1/σ_φ and K_2/σ_φ .

B. Two Hypotheses, Three Models

We now use (8) as a framework for testing two broad hypotheses that might explain the greater propensity of both young men and women to report helpful and hurtful discrimination. The first of these is greater heterogeneity in the gender-related treatment experienced by the young: large amounts of both hurtful and helpful discrimination may be more common among the young, for example, if the young find themselves in a much greater variety of jobs and work situations than their parents.

There are two main ways that this idea of "greater heterogeneity" might enter the model in (8). One is that young women could be endowed with a distribution of observed characteristics (Z 's) that generates more reports in both tails; to assess whether this kind of heterogeneity can explain the pattern of reports in our data, we simply estimate (8) as it stands, then ask whether (given the actual distribution of observed characteristics in our data) it can successfully mimic the pattern of reports across age groups in the data. In what follows, we shall refer to this ordinary ordered probit as "model 1."

Maximum likelihood coefficient estimates for model 1 are shown in column 1 of table 8 for women and column 1 of table 9 for men.²⁵ They show the following. First, as in the previous section, only a few observed characteristics (Z 's) have significant direct effects (θ 's) on the total amount of net favorable treatment faced by individual women or men. These include the percent female in one's occupation, marital status, and education; they do not include age. As expected, high values of percent female in one's occupation significantly increase (decrease) the probability of reporting helpful (hurtful) discrimination for women. Although not significant, the reverse is true for men: thus both women and men are

²⁵ A key method of assessing our models' performance in this section involves comparing the predicted pattern of reports across age groups with the unadjusted means in tables 2 and 7. For that reason we represent age by a set of four dummy variables, instead of a continuous variable, throughout this section.

Table 8
Ordered Probit Models of Gender-Induced Harm and Advantage: Women

	Model 1	Model 2	Model 3
Age 25–34	-.0982 (.78)	-.0956 (.82)	-.1236 (.99)
Age 35–44	-.0945 (.72)	-.1162 (.97)	-.1859 (1.38)
Age 45–54	-.0360 (.24)	-.0508 (.38)	-.1009 (.67)
Age 55–64	.0010 (.01)	-.0780 (.41)	-.1983 (.74)
Education:			
High school	-.0789 (.76)	-.0832 (.92)	-.0833 (.79)
Some college or university	.0821 (.62)	.0557 (.48)	.0777 (.58)
College degree	.1377 (1.09)	.1169 (1.05)	.1349 (1.05)
University degree	-.0842 (.65)	-.1149 (1.02)	-.0992 (.76)
Married	.1970 (1.99)	.1725 (1.97)	.1934 (1.97)
Separated/divorced/widowed	.0463 (.33)	.0083 (.07)	.0206 (.15)
Children less than 6 years old	.0118 (.16)	.0215 (.33)	.0141 (.19)
Tenure (ROE job)	-.0000 (.14)	-.0000 (.28)	-.0000 (.16)
Percent female in occupation	.4385 (2.57)	.4087 (2.73)	.4612 (2.66)
c_1	-1.10 (6.87)	-.9704 (6.32)	-.9680 (5.70)
c_2	1.99 (11.72)	1.71 (9.22)	1.75 (9.09)
d_2			-.0849 (.79)
d_3			-.2988 (2.66)
d_4			-.2428 (2.02)
d_5			-.6352 (2.66)
b_2		-.0649 (.87)	
b_3		-.2013 (3.17)	
b_4		-.1747 (2.48)	
b_5		-.3219 (3.61)	
Log likelihood	-705.71	-697.44	-698.13

NOTE.—1,466 observations in all three models. Province dummies were included in all three models. Absolute value of z -statistics in parentheses. ROE = Record of Employment.

more likely to feel harmed when they find themselves in a “minority” in their occupation. Further, marital status also has opposite effects on the underlying “net favorable treatment” index, Δ , for men and women, with the effect being significant for women. Compared to less than high school—the omitted category—men’s probability of reporting hurtful (helpful) discrimination is significantly higher (lower) at the two highest levels of education (college and university). The same is not true for women.

Can model 1 successfully explain the pattern of reported gender-induced harm and advantage in our data? To address this issue, the predicted pattern of reported discrimination from the model across age groups is shown in part B of table 10. These predictions give the mean predicted fraction of individuals in each age category reporting discrimination, with each age category evaluated at its own mean characteristics, Z . Thus they are directly comparable with the sample means for these variables in tables 2 and 7, reproduced for convenience in part A of table 10. Clearly, model 1 cannot successfully mimic the pattern of lower reports of both harm and advantage among younger women in our data.

A second way in which greater heterogeneity for the young could enter

Table 9
Ordered Probit Models of Gender-Induced Harm and Advantage: Men

	Model 1	Model 2	Model 3
Age 25–34	.0636 (.56)	.0575 (.53)	.0588 (.53)
Age 35–44	.0940 (.75)	.0819 (.68)	.0852 (.70)
Age 45–54	-.0172 (.12)	-.0322 (.24)	-.0420 (.29)
Age 55–64	.0389 (.22)	.0320 (.19)	.0349 (.19)
Education:			
High school	-.1799 (2.04)	-.1815 (2.13)	-.1875 (2.11)
Some college or university	-.1286 (1.13)	-.1367 (1.25)	-.1375 (1.21)
College degree	-.2341 (1.96)	-.2254 (1.96)	-.2353 (1.97)
University degree	-.2741 (2.12)	-.2940 (2.35)	-.2908 (2.23)
Married	-.0320 (.34)	-.0215 (.24)	-.0283 (.30)
Separated/divorced/widowed	-.1038 (.71)	-.1081 (.77)	-.1078 (.72)
Children less than 6 years old	.0000 (.00)	-.0042 (.07)	-.0016 (.03)
Tenure (ROE job)	-.0001 (.60)	-.0001 (.70)	-.0001 (.65)
Percent female in occupation	-.2672 (1.69)	-.2617 (1.73)	-.2702 (1.70)
c_1	-1.76 (13.14)	-1.70 (11.51)	-1.70 (11.51)
c_2	1.57 (11.89)	1.47 (10.04)	1.48 (9.76)
d_2			-.0515 (.54)
d_3			-.0117 (.12)
d_4			-.2860 (2.39)
d_5			-.1862 (1.21)
b_2		-.0358 (.61)	
b_3		-.0051 (.08)	
b_4		-.1578 (2.65)	
b_5		-.1173 (1.43)	
Log likelihood	-805.14	-800.45	-800.68

NOTE.—2,057 observations in all three models. Province dummies were included in all three models. Absolute value of z-statistics in parentheses. ROE = Record of Employment.

our model is through unobserved factors rather than observed ones, that is, through a higher level of σ_φ . To explore this possibility, let “model 2” be a variation of (8) in which:

$$\sigma_\varphi = \sigma_0 + \sum_2^5 b_k A_{k\varphi} \quad (9)$$

where the $A_{k\varphi}$ are dummies for each of the age groups in the sample. It is then easy to see how age, working through the σ_φ , can increase reports in both tails of the distribution. Again, while the relative variance terms for different age groups, that is, the b_k s, are identified, the overall “baseline” variance is not; we thus set σ_0 equal to one and interpret the b_k accordingly. To the extent the estimated variance of unobserved differential treatment declines with age, and that this model successfully mimics the pattern of reported discrimination across age groups in the raw data, we can conclude that a possible cause of the greater reports of hurtful

Table 10
Predictions of Gender-Induced Harm and Advantage by
Age Category: Alternative Models

	Age Category	Hurt	No Impact	Help
A. Sample means:				
Females	55-64	.043	.947	.011
	45-54	.088	.878	.034
	35-44	.093	.880	.027
	25-34	.123	.826	.051
	< 25	.132	.805	.064
Males	55-64	.045	.917	.038
	45-54	.042	.929	.030
	35-44	.058	.878	.063
	25-34	.061	.889	.050
	< 25	.067	.875	.058
B. Model 1: Fixed cutoffs and variance:				
Females	55-64	.085	.867	.047
	45-54	.095	.862	.043
	35-44	.106	.855	.039
	25-34	.108	.854	.038
	< 25	.112	.852	.035
Males	55-64	.051	.897	.052
	45-54	.056	.899	.045
	35-44	.047	.899	.055
	25-34	.050	.899	.050
	< 25	.055	.900	.046
C. Model 2: Variance varies with age:				
Females	55-64	.051	.935	.014
	45-54	.087	.878	.036
	35-44	.093	.880	.027
	25-34	.122	.826	.051
	< 25	.140	.799	.061
Males	55-64	.039	.921	.041
	45-54	.038	.935	.028
	35-44	.054	.884	.062
	25-34	.053	.895	.052
	< 25	.063	.883	.054
D. Model 3: Cutoffs vary symmetrically with age:				
Females	55-64	.048	.940	.012
	45-54	.089	.874	.037
	35-44	.094	.878	.028
	25-34	.121	.829	.050
	< 25	.138	.803	.059
Males	55-64	.040	.918	.042
	45-54	.038	.935	.027
	35-44	.054	.885	.062
	25-34	.053	.894	.053
	< 25	.063	.883	.054

discrimination among young women is greater unobserved heterogeneity in the labor market experiences of such women.

Coefficient estimates for model 2 are shown in column 2 of tables 8 and

9; for the most part, all the observed covariates (Z) have coefficients of similar size and significance compared to column 1 of tables 8 and 9. Concerning the variance parameters (b), we find that the estimated variance of unobserved evidence decreases as we move “up” the age categories, beginning with the omitted category (age 15–24). The standard deviation is significantly higher among the youngest age group, compared to the oldest. Further, as table 10 shows, model 2 is also much better at mimicking the patterns of reported discrimination across age groups for both men and women in our data. Overall, we take our estimates of models 1 and 2 together as mixed support for the “heterogeneity” hypothesis, because if the experiences of the young really are more heterogeneous, we would expect at least some of this to be captured by their observed characteristics in model 1. Given this mixed message, we now turn to a final hypothesis.

The final hypothesis we consider in this article involves not a difference in the “actual” labor market experiences of the young and old but a difference in how their experiences are translated into survey reports of harm or advantage. In particular, we now ask what happens when we hold σ_φ constant, but let the amount of departure from gender neutrality, K , which induces reports of hurtful or helpful discrimination, vary across age groups. Noting from (7) and (8) that an increase in K has the effect of reducing K_1 and increasing K_2 by equal amounts, this is equivalent to specifying

$$K_1 = \overline{K}_1 + \sum_2^5 d_k A_k \quad (10)$$

and

$$K_2 = \overline{K}_2 - \sum_2^5 d_k A_k. \quad (11)$$

The idea is, in a sense, that younger women (and perhaps younger men) are less “tolerant” of (or more sensitive to) departures from gender neutrality, in the sense that these are more likely to induce reports of harm or advantage, than older people. It is worth noting that this notion is, at least in principle, empirically distinguishable from the previous version of the model where unobserved heterogeneity, σ_φ , varied across age groups;²⁶ thus we estimate model 3 as well, and ask how well it can mimic actual patterns of reported discrimination in the data.

²⁶ To see that these “unobserved heterogeneity” and “differences in thresholds” models have different empirical implications, consider the effect of σ_φ and K_1 on

Coefficient estimates from model 3 are shown in column 3 of tables 8 and 9. As we might expect, the estimated cutoffs for reporting harm or advantage move farther apart as we move up the age ladder.²⁷ And, as part D of table 10 shows, model 3—like model 2—does a successful job of reproducing the pattern of increasing reports in both tails found among young women and men in the raw data. Indeed, while there are some subtle conceptual differences between models 2 and 3, both yield essentially the same maximized value of the log-likelihood function, and (because they have the same number of parameters as well) are essentially indistinguishable statistically. On the other hand, testing model 1, which has four fewer parameters, against model 3 yields a likelihood ratio (LR) test statistic for women of 15.15 with a p -value of 0.004; for men of 8.93 with a p -value of 0.063. Similar test results are obtained when model 2 is used as the unrestricted model; thus both models are preferred to the simple ordered probit specification of model 1.

Conceptually, both greater heterogeneity in the experiences of young people and less “tolerance” of departures from gender neutrality among the young could of course explain the greater tendency of young men and young women to report both hurtful and helpful discrimination in our data. In this section, however, we have shown two things. First, if it is differences in the amount of heterogeneity that matters, it must be heterogeneity that is not captured by any of the independent variables in our data set, that is, what we call heterogeneity on unobserved dimensions. Second, a model that allows for differentials in either the variance of unobserved heterogeneity or in reporting thresholds across age groups is statistically preferred for both women and men to one that does not.

V. Conclusions

Virtually all the standard, “objective” measures of discrimination computed by economists, including unadjusted or residual wage gaps, and unadjusted or residual gaps in unemployment durations, tend to be higher among older than younger women. Whether this represents true age effects or a difference across cohorts remains an open question, but it is clear that, among women in the labor market today, discrimination as we usually measure it is much more prevalent among older women. At the same time, and perhaps paradoxically, younger women are much more

the left-hand side of (8), which determines the fraction of individuals reporting hurtful discrimination. An age-related change in σ_φ changes the sensitivity of the left-hand side to the observed covariates, Z . An age-related change in K_1 does not.

²⁷ A negative d_2 , e.g., reduces the lower threshold (thus reducing reports of hurtful discrimination) while at the same time increasing the upper threshold and reducing reports of “helpful” discrimination.

likely than older women to report, in confidential surveys, that they have suffered from sex discrimination or experienced gender-induced harm in the labor market. This lack of correspondence between the two measures, in our view, should raise serious questions about whether “objective” measures really capture what most women really see as discriminatory treatment in labor markets today.

This article has attempted to resolve the apparent inconsistency between age patterns of reported and measured discrimination using a new data set on Canadian job seekers. Significant advantages of this data include a richer set of “objective” labor market outcome variables, including unemployment durations, reservation wages, and wage changes, in addition to the wage levels that are usually examined, and a richer set of perceived discrimination measures than in previous studies. In contrast to earlier work, we have information not only on women’s perceptions of gender-induced harm but also on women’s perceptions of gender-induced advantage, and on men’s perceptions of advantage and harm.

Our main findings are as follows. First, young women’s more frequent reports of gender-induced harm cannot be statistically attributed to any observable differences between them and older women, including the presence of children and degree of occupational segregation. Interestingly, we do find that minority status in an occupation increases reports of gender-induced harm among both women and men. However, as minority status is not correlated with age in our data, this does not help explain the age pattern of reports.

Second, as expected, young women’s more frequent reports of harm also cannot be attributed to a higher level of “objectively” measured discrimination on any of a large set of dimensions we can measure, including wages and unemployment. Third, the newly available information on perceptions of discrimination in our survey allows us to rule out another potential explanation of young women’s greater reported harm: a higher mean level of unmeasured discrimination against young women. This is because of a surprising and robust symmetry in the correlation of age with reports of discrimination in our data: reports of both hurtful and “helpful” discrimination fall with age in our data, and do so for both women and men.

Fourth, we identify two hypotheses that are potentially consistent with all the patterns we observe. One of these—greater dispersion in the labor market experiences of young workers—receives mixed support, in the sense that it is consistent with our data only if none of these differences in dispersion can be captured by dispersion in the observed covariates in our data. Alternatively, our results could be explained by a very particular kind of age difference in reporting behavior: young workers of both sexes may simply be more willing than older workers to interpret departures in either direction from gender-neutral treatment as causally related to their

gender. Interestingly, this difference in reporting behavior takes exactly the form we would expect if, in conducting implicit statistical “tests” about whether they face discrimination based on the information available, younger persons used a less stringent significance criterion.

In future work, other demographic patterns in reported discrimination may of course also be identified and analyzed.²⁸ At the same time, additional explanations of age and other demographic patterns in reports are likely to emerge. What seems clear from the current analysis, however, is that all such hypotheses will need to somehow be consistent with the surprising symmetry in responses of helpful and hurtful discrimination to age that emerges so strongly from our data here. This stronger perceived effect—in both directions—of gender on young peoples’ labor market outcomes must somehow figure into those new stories.

Finally, it may be worth noting that even if we adopt the pure “reporting differences” interpretation of our results, those results may still have important policy implications. For example, if lower “tolerance” of nongender-neutral treatment remains a permanent attribute of today’s cohort of young workers, designers of future antidiscrimination policies face a dilemma: while the young are less tolerant of discrimination, they also appear to be less tolerant of reverse discrimination. Creating policies that address the former problem without creating perceptions of the latter may thus become increasingly difficult in coming decades.

Appendix A

Table A1
“Objective” Wage Discrimination, by Age Group: Estimates from the 1985 Panel Study of Income Dynamics

Age	Measured Discrimination: Log Wages		
	Unadjusted Gap (1)	Adjusted Gap (\hat{D}^1) (2)	Adjusted Gap (\hat{D}^2) (3)
55–64	.519	.465	.499
45–54	.594	.455	.420
35–44	.465	.340	.334
25–34	.326	.233	.258
18–24	.219	.258	.202

²⁸ A prime candidate here is education effects. As noted earlier, education effects are stronger for men than women but exhibit the same “symmetry” property as was identified here for age. Patterns of reported discrimination against various ethnic and racial groups also deserve analysis.

Appendix B

Table B1
Probit Coefficients for Gender-Induced Harm or Advantage

	Women	Men
Age	-.0116 (2.25)	-.0007 (.16)
Education:		
High school	-.0673 (.57)	.0025 (.02)
Some college or university	-.0703 (.47)	.3122 (2.55)
College degree	.0929 (.67)	.1848 (1.42)
University degree	.2679 (1.93)	.1890 (1.35)
Married	-.1718 (1.61)	-.0995 (.96)
Separated/divorced/widowed	.1733 (1.15)	-.1019 (.62)
Children less than 6 years old	-.0444 (.54)	.0355 (.53)
Tenure (ROE job)	-.0000 (.26)	-.0001 (.80)
Percent female in occupation	-.3631 (1.90)	.2025 (1.17)

NOTE.—Absolute value of z-statistics in parentheses. ROE = Record of Employment. In addition to the control variables above, the probit regressions include provincial dummy variables.

References

- Barbezat, A., and Hughes, J. W. "Sex Discrimination in Labor Markets: The Role of Statistical Evidence: Comment." *American Economic Review* 80 (March 1990): 277–86.
- Hampton, Mary B., and Heywood, John S. "Do Workers Accurately Perceive Gender Discrimination?" *Industrial and Labor Relations Review* 47 (October 1993): 36–49.
- . "Racial Earnings Differences and Worker Perceptions: The Case of Physicians." Photocopied. La Crosse and Milwaukee: University of Wisconsin, October 1996.
- Heywood, John S. "Race Discrimination and Union Voice." *Industrial Relations* 31 (Fall 1992): 500–8.
- Johnson, George, and Solon, Gary. "Estimates of the Direct Effects of Comparable Worth Policy." *American Economic Review* 76 (December 1986): 1117–25.
- Johnson, Richard W., and Neumark, David. "Age Discrimination, Job Separations, and Employment Status of Older Workers: Evidence from Self-Reports." *Journal of Human Resources* 32 (Fall 1997): 779–811.
- Kuhn, Peter. "Sex Discrimination in Labor Markets: The Role of Statistical Evidence." *American Economic Review* 77 (September 1987): 567–83.
- . "Sex Discrimination in Labor Markets: The Role of Statistical Evidence: Reply." *American Economic Review* 80 (March 1990): 290–97.
- . "Demographic Groups and Personnel Policy." *Labour Economics* 1 (June 1993): 49–70.
- Laband, David N., and Lentz, Bernard F. "Is There Sex Discrimination in the Legal Profession? Evidence on Tangible and Intangible Margins." *Journal of Human Resources* 28 (Spring 1993): 230–58.
- . "The Effects of Sexual Harassment on Job Satisfaction, Earnings, and Turnover among Female Lawyers." *Industrial and Labor Relations Review* 51 (July 1998): 594–607.
- Oaxaca, Ronald. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review* 14 (October 1973): 693–709.