2023

Modeling the Performance-based Compensation of MLB Catchers

William DeForest

Claremont McKenna College

Modeling the Performance-based Compensation of MLB Catchers

submitted to
Professor Ricardo Fernholz

by
William DeForest

for
Senior Thesis
Fall 2022
5 December 2022

Table of Contents

Acknowledgments

I would like to thank Professor Ricardo Fernholz and Professor Andrew Finley for their guidance over the past semester. I wouldn't have been able to complete this thesis without your mentorship and constructive advice. In addition, I would like to thank Kellen Sundin who inspired my love for catching and helped spark the idea for this project. I would also like to thank my family and friends for their support for helping me complete this work.

Abstract

How professional baseball players are compensated for their on-field performance and contribution to winning baseball games has been studied many times by baseball analysts, yet no real attempt has been made to focus solely on the catcher position. As one of the most demanding and specialized positions in the game, a talented catcher is vital to a team's success. This paper attempts to utilize Defensive Runs Saved (DRS) and Offensive Runs Above Average (OFF), leading statistics that use Statcast data to measure defensive and offensive productivity, to determine which area of catcher defense is the greatest differentiator between elite defensive catchers and average defensive catchers and to determine how teams are compensating catchers for their on-field performance.

Analyzing data collected on catchers who played in the MLB between 2016-2022, I find that strike zone runs saved above average (rSZ), which measures the pitch receiving ability of a catcher, has the highest standard deviation (4.65 runs saved) of all the DRS components by a statistically significant margin. This implies that teams should prioritize signing, and should be willing to pay more for, catchers who are elite receivers because this ability will save the team more runs over the course of the season than if they sign a catcher who is elite in another defensive component. To determine how teams compensate catchers for their on-field performance, I regress DRS, OFF, and their various components on annual salary while controlling for contract type and year. I find that a one run increase in DRS causes a 1.2% increase in salary while a one run increase in OFF causes a 1.1% increase in salary. This indicates that teams are efficient with their compensation of catchers and pay them nearly evenly for their overall defensive and offensive contributions.

1. Introduction

Major League Baseball (MLB) is a massive organization projected to generate close to $11 billion in revenue in 2022. Every year, the 30 teams in the MLB compete fiercely to be crowned the World Series Champion and win tens of millions of dollars in prize money, and even more in media deals and fan revenue. To this end, MLB teams collectively spend billions of dollars on players' salaries annually. With so much at stake, it is to be expected that teams and owners want to know how they can win more, while spending less. This question can be asked about each of the nine positions on the field and answering it for the catcher position, arguably one of the most important to a team's success, is the focus of this study.

Essential to this study are the detailed statistics, dubbed "sabermetrics," that are at the heart of every transaction completed in the MLB today. They provide the most concrete evidence for how a player will perform in the future. However, despite much discussion, no real consensus has been reached and researchers and analysts are constantly searching for new, more accurate ways to measure how good a player is and how much they are worth. This study builds on prior research by Hakes and Sauer (2006), Brown, Link and Rubin (2015), and Pollack (2017) linking on-field performance to compensation but is unique in focusing solely on the catcher position.

The catcher plays a crucial role on both offense and defense. On offense, they are one of the nine players who form the batting lineup and take turns facing the pitcher. On defense, catchers are active every pitch - receiving, throwing, calling pitches, and occasionally blocking errant balls thrown by the pitcher. As one of the most demanding and specialized positions in the game, finding a talented catcher is vital to a team's success. Luckily, thanks to recent advances in technology, this search is being aided by extremely detailed performance data and statistics. By 2015, all MLB ballparks instituted Statcast, a system designed to track the baseball and every player on the field simultaneously. There has been a considerable amount of work done to harness this newly available data using various measures; key among these is the introduction of the Defensive Runs Saved (DRS) sabermetric statistic. Created by *The Fielding Bible*, DRS uses Statcast data as inputs and calculates how many runs better or worse that player has produced relative to the average player at his position. It breaks down the defensive side of catching into six measurable components and has given teams, analysts, and fans alike a more accurate way to judge a catcher's defensive productivity.

The DRS statistic is crucial to the two main goals of this paper: to determine which area of catcher defense, as measured by DRS, is the greatest differentiator between elite defensive catchers and average defensive catchers and to determine how teams are compensating catchers for their defensive and offensive productivity.

With these goals in mind, the two hypotheses driving this paper are:

1. Pitch receiving is the greatest source of variation in defensive productivity at the catcher position.
2. Offensive productivity is compensated more than defensive productivity.

The primary reason behind Hypothesis One is the frequency that the pitch-receiving skill is tested. Every single pitch of the game is an opportunity for catchers to give their team an advantage - even if that advantage is small. Even if there is just a small difference in receiving ability between two catchers, the difference will be magnified massively when compounded over the thousands of pitches catchers receive each season. As for Hypothesis Two, it is motivated by the idea that offensive productivity is more valued by MLB teams because it drives more fan attendance than defense. However, a study by Ehrlich and Potter (2020) found that fans actually have no preference for teams with good offense over teams with good pitching or defense. Their results suggest that team decision-makers act irrationally by paying more for offense than they do defense and that defensive and pitching wins should be valued at the same rate as offensive wins on the free agent market. This study will add to these previous findings by focusing solely on the catcher position and by breaking down offense and defense into specific skills.

This study uses DRS as a proxy for a catcher's defensive productivity and analyzes the standard deviations of its various components to determine which areas of catcher defense drive the difference in productivity between two catchers. Data is sourced from the online baseball statistics database fangraphs.com and includes all catchers who caught more than 90 innings in a season in the MLB between 2016 and 2022. The resulting dataset has 427 observations, each of which has the catcher's DRS and DRS component statistics. Analysis of these statistics finds that strike zone runs saved, at 4.65 runs saved, has the highest standard deviation of all the DRS components by a large margin. This supports the hypothesis that pitch receiving is the area of catcher defense with the most variation in productivity. This means that teams should prioritize signing, and should be willing to pay more for, catchers who are elite receivers because this ability will save the team more runs over the course of the season than if they sign a catcher who is elite in another defensive component. This also suggests that younger catchers should focus their practice time on receiving, as becoming a better receiver will have a bigger impact on their overall defensive ability than becoming better at any other individual area of catcher defense.

To test Hypothesis Two, a measure of offensive productivity and a measure of compensation are needed in addition to DRS. I use the offensive runs above average (OFF) statistic because it is also measured in terms of runs above average and is therefore easily comparable to DRS. As for compensation, an ideal measure of player compensation would include details such as signing bonus and other incentives. However, this information is not publicly available. Instead, this study uses the player's salary as of the first game of the season as a proxy for compensation. In addition to salary, the number of years of Major League Service (MLS) accumulated, team, and player age are also sourced from Cot's Baseball Contracts, a Baseball Prospectus database for MLB salary and payroll details.

MLS time is a particularly important control variable because it determines which type of contract the player is eligible for. The MLB breaks player compensation down into three stages: team-mandated, arbitration, and free agency. Each stage is tied to how many years of service a player has, and at each stage players become eligible for substantial increases in salary. However, due to the lack of heterogeneity in salaries of players on team mandated contracts, I choose to exclude this group of players from the analysis.

Linear regression models are used to test the second hypothesis that offensive productivity is compensated more than defensive productivity. Since DRS and OFF are easily comparable, identifying whether the results support my hypothesis that offensive productivity is over-valued by teams is as simple as comparing the coefficients for DRS and OFF. If teams are acting rationally and the contract market is efficient, we would expect teams to pay the same amount for a run gained on offense as a run saved on defense and therefore for the coefficients on DRS and OFF to be equal. My hypothesis is that the coefficient for OFF will be higher than the coefficient for DRS (that offense is overvalued, and defense undervalued) because of the misperception that fans enjoy watching offense more than defense.

However, this hypothesis is rejected by the findings from the regression results. Results show that teams in fact pay slightly more for a one run increase in DRS than they do for a one run increase in OFF. Specifically, a one run increase in DRS causes a 1.2% increase in salary while a one run increase in OFF causes a 1.1% increase in salary. This indicates that teams are efficient with their compensation of catchers from an on-field performance standpoint. If anything, teams pay slightly more for defensive productivity than they do offensive productivity. In fact, there is evidence that teams are aware of the importance of the pitch receiving component of catcher defense because after breaking down DRS and OFF into their respective components, results show that a one run increase in rSZ causes a 2.9% increase in salary. This is greater than the estimated 2.12% salary increase from an additional one run of batting. These findings imply that teams accurately compensate catchers and do not operate under the misperception that offensive productivity is more valuable because it is more exciting for fans.

The findings of this paper differ from prior literature such as Maurice (2010) and Ehrlich and Potter (2020) who found that teams pay more for offensive performance than they do for defensive performance. However, these prior studies look at all field positions, rather than at the catcher position specifically. This is an important difference because of the extremely specialized nature and importance of the catcher position on defense. Other than perhaps the pitcher, the catcher has the greatest impact on the game when a team is on defense. As such, it is reasonable that teams prioritize defensive productivity at the catcher position when making roster and salary decisions.

While this study is limited by the lack of detailed compensation information and a relatively small sample size since it analyzes only MLB catchers, it has many interesting potential extensions such as repeating the analysis for other positions. As such, this study offers a repeatable framework for analyzing the compensation of players at each position in the MLB based on their offensive and defensive productivity.

The remainder of the paper proceeds as follows. I begin by providing context about sabermetrics and the importance of the catcher position. I then introduce my hypotheses and their motivations before exploring the MLB contract structures in more detail. The following section describes my methodology including the data, variables, and models used in testing the hypotheses. Next, I discuss the findings from exploratory data analysis including descriptive statistics and correlations. Lastly, I present the main results and discuss their implications. The appendix includes the discussion of two robustness checks, the results from two alternative model specifications, and rankings of the best performing individual catchers and teams (as a group of catchers) over the period of analysis.

2. Background

A. Sabermetric Revolution

In the late 1970s, the world of baseball experienced a statistical renaissance which has led baseball to become one of the biggest data-driven sports worldwide. This revolution began when Dick Cramer, Bill James, and Pete Palmer co-founded the Society of American Baseball Research's (SABR) Statistical Analysis Committee. James became the figurehead of the movement after publishing his extremely popular *Baseball Abstract* books and introducing the term "sabermetrics" (Mizels 2022). Sabermetrics, whose name is taken in part from the Society of American Baseball Research's acronym "SABR", is defined by James as "the search for objective knowledge about baseball" (Kelly 2019). To James and other "sabermetricians" this practice involves utilizing detailed statistics to make empirically supported decisions regarding which players a team should pursue and what value a player has to a baseball team (Beneventano 2012).

This approach was markedly different from the traditional method of evaluating players based largely on qualitative factors such as height, weight, and attitude and a few relatively simple statistics such as batting average (Beneventano 2012). James and other sabermetricians believed that the traditional statistics such as batting average (BA), runs-batted-in (RBI), and earned-run-average (ERA) did a poor job of measuring a player's or team's performance. In his 1979 Baseball Abstract James wrote "it is startling… how much confusion there is regarding how a hitter or team should be measured." As an example, James often cited how the batting average statistic, which is used to quantify offensive production, excludes walks, hit by pitches, sacrifices, and catcher's interferences (Mizels 2022). These can be valuable plate appearances for a team and

James pointed out that ranking teams by BA resulted in teams with higher run totals being ranked below teams with fewer runs scored. Given that the goal of a team's offense is to score as many runs as possible, this is evidence that batting average is not a particularly accurate measure of offensive performance. To remedy these inaccuracies, Sabermetricians developed new statistics, dubbed "sabermetrics," by using conventional statistics in carefully chosen combinations to calculate measures thought to gauge a player's value or relative worth more accurately (Beneventano 2012).

The sabermetric movement faced initial resistance by professional scouts who saw it as a disturbance of baseball tradition. However, these scouts were biased as their jobs were predicated on their capacity to judge players' qualitative and intangible abilities, and the introduction of sabermetrics threatened to make the scouts' judgment skill less important. As a result of the pushback from scouts, it wasn't until the early 2000s and the publishing of *Moneyball* (2003), a story written by Michael Lewis about how Billy Bean, general manager of the Oakland Athletics, used sabermetrics to take advantage of trade market inefficiencies and acquire players at a fraction of their true value, that sabermetrics really caught hold. *Moneyball* was a national sensation and today every MLB team has made sabermetrics an integral part of their decision-making process (Avidon 2022).

B. Importance of Catchers

Over the course of a 162-game season, no other player, other than perhaps an ace pitcher, has a greater impact on team performance than the catcher. It is the most demanding position in the game - both physically and mentally (Garro 2019).

The daily grind of the catcher position is unparalleled. A starting catcher crouches or kneels behind home plate for nine innings straight in likely four to five games per week. In comparison, the best starting pitchers only pitch once every five games and even then, only pitch on average between five and seven innings. Catchers make the most throws of any position each game because every time a pitcher makes a pitch, the catcher must throw the ball back. Of course, the throw back is at a lower intensity, but the sheer volume is astounding. Every time a ground ball is hit when no one is on base, the catcher runs - fully geared - to back up first base. Occasionally even, when there is a tag play at home plate, the catcher is run or slid into by the sprinting runner trying to score. No other position on the field demands as much from the body day in and day out as catching.

The mental side of the game is just as taxing at the catcher position. One of their most important jobs is to call pitches. To do this well requires the analysis of a myriad of information including batter scouting reports, pitcher abilities and tendencies, game situation, where the batter is standing in the batter's box, weather conditions and much more. This aspect of the game has so far proven difficult to quantify as there is little observable data available to measure pitch-calling ability. Furthermore, catching major league pitching, especially at a high level, is extremely difficult and relies upon exceptional hand-eye coordination and very fine motor movements. To perform their job well, catchers are required to maintain a laser-focus throughout the ballgame. A final important responsibility of the catcher position is to know their pitchers almost better than they know themselves. If a pitcher is struggling with command or is giving up a lot of hits, the catcher needs to know how best to calm their teammate down and get them locked back in. This makes managing the relationships with the pitching staff an important and delicate task. Catchers need to have a high social I.Q. in addition to a high baseball I.Q.

On top of all their defensive responsibilities, catchers need to hit too. Like all the other position players, catchers hit once every nine batters. While historically catchers haven't been as good hitters as the other positions and generally, expectations for offensive performance are lower than at other positions, catchers are still an integral part of a team's offense.

Being such a difficult and specialized position, one might expect good catchers to be a hot commodity. Curiously however, this is not reflected in the compensation of Major League Baseball catchers. In fact, catcher is the lowest paid position in baseball ("MLB Positional Payrolls"). At $1,604,864, they have a lower average salary than even relief pitchers who, while still being important to a successful ballclub, play only a fraction of the number of innings catchers play.

Recent advances in technology have transformed how MLB catcher performance is being evaluated today. In 2008, Sportvision's PITCHf/x system was operationalized in all 30 MLB ballparks (Healey 2017). This system applied computer vision techniques to video obtained from cameras located around the stadium to estimate the 3-D path of pitched balls. This system opened a new realm of sabermetric possibilities regarding the speed, movement, and location of pitches. It provided objective pitch location data which finally allowed the catcher's pitch receiving ability to be quantified. Many coaches, scouts, players, and statisticians were quite surprised by the initial results for how impactful pitch receiving is in baseball. In 2011, Mike Fast wrote the first substantial study that attempted to use the new data to judge catchers' receiving abilities and found that this skill alone could gain or cost their team up to two wins in a season (Fast 2011). These findings re-shaped how teams think about catcher defense and fostered a new, receiving-focused style of catching (Kuty 2022).

Since 2008, ball-tracking technology has continued to improve. In 2015, Statcast, a system designed to track the ball and every player on the field, was functional in all MLB ballparks. It replaced the PITCHf/x system and brought improvements to pitch-tracking accuracy, particularly with measuring the spin of the ball, as well as the added capability of being able to track players.

3. Hypotheses Development

A. *Hypothesis One: Pitch receiving is the greatest source of variation in defensive productivity for catchers*

There has been a considerable amount of work done attempting to harness the newly available data to better understand what makes a catcher a better defensive player and how valuable those defensive abilities are to a team. Neil Weinberg of the popular baseball statistics database "fangraphs.com" breaks down catcher defense into five categories (Weinberg "What do we know" 2014):

1. *Normal fielding*

Normal fielding includes fielding bunts, catching pop-ups, and tagging out runners trying to score. While these "traditional" plays are straightforward to judge when they do occur, they just don't occur very frequently. For this reason, there isn't a significant amount of research being conducted in this area of catcher defense.

2. Pitch receiving

In contrast to normal fielding, pitch receiving (or framing) is a buzzing area of research for baseball analysts right now. In addition to Mike Fast's prominent 2011 paper, there are dozens of studies attempting to measure the pitch receiving ability. In one of the most recent papers on the topic, DaSilva describes two potential receiving metrics (DaSilva 2021). Framing runs above average (FRAA), weighs each pitch with an expected run value, derived from the difference in expected runs of each given ball-strike count had the umpire made the correct call and the count following a stolen or lost strike. Using this metric, catchers only gain or lose runs if they steal or lose a strike call. Catcher strikes above average (CSAA) does not value a catcher's framing performance based on runs, but rather on how difficult a pitch is to frame for a strike. Each pitch is given a called strike probability. Using a K-nearest neighbor regression model, which considers the pitch location, pitch type, batter handedness and ball-strike count, probabilities are predicted for each pitch to be called as a strike by the umpire. Every called strike a catcher frames, they gain the difference of the corresponding probability from one; for every lost strike, they lose the difference to zero.

3. Blocking

The third area of catcher defense that Weinberg identifies is blocking. Blocking occurs when a pitcher throws a ball that bounces in the dirt in front of the catcher. It is an important skill because if the ball gets past the catcher or bounces off him and rolls far enough away, a baserunner can advance to the next base. One contemporary metric which measures the blocking ability of catchers is the passed pitch statistic. The author

Brian Koprivica broke up where pitches cross the plate (or land if the ball is in the dirt) into buckets and determined the probability of a passed pitch occurring in each of those locations (Koprivica 2011).

4. *Game calling*

As mentioned briefly above, researchers are yet to have found a way to objectively measure pitch-calling ability. Part of the issue is that it is difficult to attribute a particular pitch call to one person. Catchers are the ones putting down the sign, but the pitcher can "shake off" the pitch until the catcher puts down the sign of the pitch that they want to throw. In addition, the extent to which the coaching and scouting staff has input on pitch-calling differs for every team. Some provide the catcher and pitcher with detailed reports on every batter on the opposing team and suggestions on which pitches to throw to get him out. Others however, usually the teams with veteran catchers who have played professional baseball for a long time, leave it up to the catcher to make the game plan. While there is much speculation, there is little verifiable research that has been done in this area. One 2014 piece by Ben Lindbergh discusses Yadier Molina's (a veteran, extremely successful catcher) pitch-calling process. In the article, Lindbergh indicates that the spread of game calling performance over the course of a season is between -10 and +10 runs, but the data driving that estimate isn't publicly available (Lindbergh 2014). As a result, game calling is still a mysterious category of catcher defense, but one that could quite possibly have the largest impact on the game from the catcher position.

5. Controlling the running game

The fifth and final area of catcher defense is controlling the running game. This involves preventing stolen bases by throwing out potential base stealers and occasionally trying to pick off runners getting too far away from the base. It is the part of catching that most casual fans likely associate with catching the most because it is an exciting play and one of the only times when the catcher is the focal point of the action. A strong and accurate arm might scare off runners and allow catchers to rack up caught stealings, but recent evidence shows that a lot of what happens on the base paths is conditional on the pitcher's time to deliver the pitch and ability to hold runners. For example, in a 2008 paper, Loughin and Bargen found evidence that pitchers have greater potential to affect stolen-base attempts and successes than catchers (Loughin and Bargen 2008). So, while throwing stealing runners out is one of the most glorified parts of catching, it doesn't add much to a catcher's defensive ability unless they are truly exceptional throwers.

With these five areas of catcher defense and my overarching goal of determining what aspects of the position drive compensation for catchers in mind, I believe that the Defensive runs saved (DRS) statistic is the best metric for measuring catcher defensive ability in this paper. DRS is a statistic calculated by *The Fielding Bible*, an organization dedicated to baseball's defensive analytics, that rates individual players as above or below average on defense. The statistic uses Sports Info Solutions data as inputs and measures how many runs better or worse that player has been relative to the average player at his position. A DRS of zero is league-average, so a positive DRS value denotes above-average performance, and a negative DRS value denotes below-average performance. The units being in terms of runs saved makes this statistic easily comparable to popular offensive statistics that are also measured in runs such as offensive runs above average (OFF) (Slowinski 2010). This makes DRS particularly well-suited for this study as we will be able to directly compare defensive and offensive ability. DRS captures a catcher's total defensive value and breaks this down into similar categories to those proposed by Weinberg. These include:

1. PART - positioning, air balls, range and throwing which essentially captures the "normal fielding" area of catcher defense
2. Bunt runs saved - fielding bunts which also falls under the "normal fielding" area
3. Good fielding play - for catchers, this is primarily the "blocking" ability of the catcher
4. Stolen base runs saved - measures the "controlling the running game" area
5. Adjusted earned runs saved - this accounts for the quality of the pitching staff the catcher works with as well as the "game calling" area
6. Strike zone runs saved - measures the pitch receiving ability of the catcher

The basic premise behind the DRS statistic is to approximate the proportion of the league's fielders who would have successfully made any given play, then to appropriately reward or penalize each fielder for his efforts (Dewan 2020). *The Fielding Bible* calls this their "plus/minus system." A play is considered "made" if a fielder gets a putout or assist on a batted ball and considered "not made" if he doesn't get a putout or assist. If the play is made, the fielder gets a positive credit and if the play is not made, he gets a negative credit. The size of the credit is directly proportional to how often that specific play is made by players at the same position across the league. The system uses Sports Info

Solutions data on location, velocity, and trajectory to group batted balls into specific plays.

For a catcher-specific example, let's say a batter hits a pop-up in foul territory. This play falls into the PART category of DRS and "normal fielding" area of catcher defense. The pop-up has a certain velocity and trajectory which the plus/minus system determines is converted into an out by the catcher 70% of the time and is not converted into an out 30% of the time. If the catcher catches the pop-up and converts it into an out, he is awarded a credit of +0.3 points. If he doesn't make the play, the catcher receives a penalty of -0.7 points. Since making plays that save doubles or triples are more valuable than making plays that save singles, an Enhanced plus/minus is also calculated to estimate how many bases a fielder saved or cost his team on a given play. Continuing the example, if 1.5 is the average number of bases a batter gains when the pop-up is not caught, then we can say that the catcher saved 1.5*.3 = 0.5 bases (Enhanced plus/minus points) on the play. These points are then converted into runs saved using the run-expectancy 24 matrix (RE24).

The RE24 matrix is based on the 24 different possibilities for baserunner/outs situations in baseball. There can be either 0,1, or 2 outs and there can be either no baserunners, just a runner on first, just a runner on second, just a runner on third, runners on first and second, runners on first and third, or runners on first, second, and third. Three outs possibilities and eight base runner possibilities combine for twenty-four total possible situations (Weinberg "RE24" 2014). There is also a "run environment" because the amount of run-scoring changes based on the time and league the teams are playing in. Run environment is calculated by taking the average team runs per game in a given league during a given year. Given a “run-environment” and base-out state, there is a specific RE24 matrix which gives the average number of expected runs per inning. To calculate the RE24 of a given play, simply take the run expectancy of the result of the play, subtract the run expectancy of the starting state, and add in any runs scored during the play (Pemstein 2016). Thus, the calculation is RE24 = RE beginning state - RE end state - runs scored.

Continuing our example from above, let’s say there is nobody on base and nobody out when the pop-up is hit. Using the RE24 matrix created by Neil Weinberg of fangraphs.com (Figure 1) below, we can see that the run expectancy of the beginning state is 0.461 runs. If the catcher makes the play and catches the pop-up, the end-state run expectancy decreases to 0.243 runs and the catcher is credited with saving the difference of 0.461 - 0.243 = 0.218 runs. If, however, the catcher misses the pop-up, the end-state run expectancy increases to 0.831 runs. In this case, the catcher is again credited with the change in run expectancy, but this credit is now negative with 0.461 - 0.831 = -0.37 runs.

Runners	0 Outs	1 Out	2 Outs
Empty	0.461	0.243	0.095
1 _ _	0.831	0.489	0.214
_ 2 _	1.068	0.644	0.305
1 2 _	1.373	0.908	0.343
_ _ 3	1.426	0.865	0.413
1 _ 3	1.798	1.140	0.471
_ 2 3	1.920	1.352	0.570
1 2 3	2.282	1.520	0.736

Figure 1: Run Expectancy 24 Matrix
This figure shows the average number of expected runs scored per inning given the current number of outs and placement of baserunners (Weinberg “RE24” 2014).

The final step to reconcile Enhanced plus/minus points and runs is to use the RE24 matrix to calculate the average number of runs that each Enhanced plus/minus point is worth for each specific position. This involves finding the average run expectancy change of this specific play across the entire season. *The Fielding Bible* calls this the "run factor". Of course, the run factor of a pop-up will be different from the run factor of another play such as throwing out a runner attempting to steal second. Therefore, DRS is broken up into different components because different areas of catcher defense have different run factors. So, to finally get us to the runs saved unit, the run factor specific to each component of DRS is used to multiply the Enhanced plus/minus points gained or lost on those types of plays by a given catcher over the course of the season.

While there are other statistics such as ultimate zone rating (UZR) that are similar to DRS, I have chosen to use DRS due to the availability of data and because UZR does not include bunt runs saved in its calculations for catchers (Dewan and Jedlovec 2009).

My first hypothesis is that pitch receiving ability, as measured by the strike zone runs saved (rSZ) component of DRS, is the most important driver of the difference in defensive performance between MLB catchers. I believe that while the other five components of DRS can be important as well, the sheer volume (every single pitch) at which receiving occurs makes it the greatest source of variation. I use the standard deviation of each component of DRS to define "importance." The components of DRS that have the highest standard deviation are those which are most responsible for the difference in defensive ability between two catchers. If, for example, five out of the six components of DRS did not vary at all between catchers, then the sixth and final component would be the only determinant of what makes one catcher better than another. Since DRS and its components are set to have a league average of 0, their means are all 0. As a result, mean cannot be used as a measure of importance.

The primary reason behind my hypothesis is the frequency that the pitch-receiving skill is tested. Every single pitch of the game is an opportunity for catchers to give their team an advantage - even if that advantage is small. The difference between a 2-1 (two balls, one strike) count and 1-2 (one ball, two strikes) count may not seem significant at first, but research by Ben Clemens on fangraphs.com shows that the weighted on base average (wOBA) of hitters in those counts are 0.361 and 0.228 respectively (Clemens 2020). That means that if a catcher can get a borderline 1-1 pitch called a strike instead of a ball, they improve their team's chance of not letting the batter reach base by 13.3%. There are similar consequences of getting a pitch called a strike in other counts as well, and when you consider that each team throws around 140 pitches per game on average, the potential impact of a catcher's receiving ability quickly adds up. Not only do more opportunities to meaningfully impact the game lead to more runs being saved overall, they also allow for a wider skill gap to develop. Even if there is just a small difference in receiving ability between two catchers, the difference will be magnified massively when compounded over the thousands of pitches catchers receive each season. As a result, while the other components of catcher defense are important too, I believe that the volume at which receiving occurs will make it the most important.

B. *Hypothesis Two: Offensive productivity is compensated more than defensive productivity*

My second hypothesis has to do with determining the specific factors that drive how much catchers are paid by their teams. Statistics are at the heart of every transaction in the MLB today because they provide the most concrete evidence for how a player will perform in the future. However, despite much discussion, no real consensus has been reached on the best methods for determining how much a player should be compensated. In fact, there is a history of market inefficiencies in the game of baseball that provide examples for times when teams realized that they weren't accurately valuing players.

One such time occurred in the early 2000s and has since been dubbed the "Moneyball" period. In one of several studies conducted looking at the player market during this time, Hakes and Sauer (2006) tested Michael Lewis' central claim in *Moneyball* (2003) that hitters' salaries did not accurately reflect the contribution of various batting skills to winning games. Specifically, they use linear regression analysis to confirm that on-base percentage (OBP) is a more powerful indicator of how much a batter contributes to winning games than slugging percentage (SLG). An efficient labor market for baseball players would, all other factors held constant, reward OBP and SLG in the same proportions that those statistics contribute to winning. However, the coefficient for SLG on the income of a player is considerably larger than the coefficient for OBP, which is the reverse of their importance to team success. Through their analysis, they provide support for Michael Lewis' claim that the valuation of skills in the market for baseball players was grossly inefficient (Hakes 2006). Hakes and Sauer follow-up their 2006 paper by extending the sample both backward and forward in time, seeking to determine how long the pricing anomaly existed, and whether the recent attenuation in the anomaly is robust to new observations. They find that the pricing anomaly extends well before the period described in Moneyball, and that with some important caveats, the market correction in the post-Moneyball period persists (Hakes 2007). The work of Hakes and Sauer is supported in a study by Brown, Link and Rubin (2015) where they test the hypothesis that in a competitive market, other teams will increase the weight given on-base percentage in the reward structure for their players. Their results show that in the post-Moneyball era, MLB teams did indeed reward players more for newer, more statistically driven performance measures (like OBP) than they did in the pre-Moneyball era. These studies show that MLB teams became aware of the power of sabermetrics in the early 2000s and began to adopt strategies based on sabermetric data.

More recently, Pollack (2017) analyzed team-level data to reveal the most important determinants of run scoring and run prevention, respectively. He then utilized models of player contract value, controlling for player-specific variables and environmental factors, to determine what factors are most significantly rewarded on the free agent market. However, he doesn't attempt to address position players' defensive contributions to run prevention effectively. In addition, Maurice (2010) compared a player's marginal revenue product to the terms of their contract and found that in 2009, players who played above average defense were consistently undervalued. The paper derives marginal revenue product by a model which attempts to estimate how a position player's individual statistics affected team winning percentage and how that winning percentage affected team revenue. However, the individual statistic used to estimate the player's impact on offense is runs scored. According to sabermetricians, this is a crude measure for estimating how much a player contributes on offense because it depends significantly on the ability of the hitters who bat behind the player to hit him in. This paper seeks to improve upon this study by utilizing higher quality statistics to estimate a player's on-field performance.

Of course, however, there is more impacting a baseball player's compensation than just their on-field performance. While how much a player contributes to winning is an important factor, players are also paid for their entertainment value. Fan attendance is an important source of revenue for MLB teams. It helps drive a better bottom line and gives the home team a psychological edge in-game which manifests itself as the "home-field advantage" phenomenon (Smith and Groetzinger 2010). Ticketing and matchday income represent about half of the total revenue a team makes in a year (Williams 2006) and a 2010 study by Smith and Groetzinger found that a one standard deviation increase in fan attendance leads to an additional half a run and a 4% increase in the likelihood of a home win (Smith and Groetzinger 2010). Furthermore, according to a 2018 paper by John Bradbury, fans prefer to allocate their entertainment dollars to winning teams (Bradbury 2018). As a result, owners and general managers have several incentives to pay players who drive fan attendance.

As for catchers and the scope of this paper, I hypothesize that, despite a run contributed on offense theoretically being just as important to winning as a run contributed on defense, offense is more valued by MLB teams because it drives more of the fan attendance than defense. This is not a novel idea and is motivated by a few key studies. First, Domazlicky and Kerr (1990) found that the introduction of the designated hitter position boosted both offense and attendance in the American League. Similarly, Tainsky and Winfree (2010) found that attendance also increased during the heightened offensive production of the 1990s steroid era. Both studies suggest that baseball fans not only prefer winning teams, but also strong offensive teams. However, Ehrlich and Potter (2020) refute this hypothesis and found that fans actually have no preference for teams with good offense over teams with good pitching or defense. Their results suggest that team decision-makers act irrationally by paying more for offense than they do defense and that defensive and pitching wins should be valued at the same rate as offensive wins on the free agent market. My study will add to this paper's findings by going into detail on the catcher position specifically and by breaking down offense and defense into their specific components. I believe that today, the perception that fans enjoy watching offense more than defense is causing MLB teams to over-value offensive productivity and under-value defensive productivity. If this is the case, teams will pay players more for a run created on offense than they do for a run saved on defense. This would be irrational since a run is a run and counts the same towards the final score whether it was scored on offense or prevented on defense.

C. *Contract Type*

The type of contract is an important covariate that must be considered and controlled for to adequately test these hypotheses. The MLB has a very specific and complicated structure for compensating its ballplayers. This structure is explicated in the MLB's Collective Bargaining Agreement (CBA) with the Major League Baseball Players Association (MLBPA) which, in addition to contracts, contains all the guidelines governing Major League Baseball except for those in the rule book about physically playing the game of baseball (Wasserman 2013). Each CBA has a term of approximately five years and the scope of this project spans two agreements: the 2017-2021 CBA and the 2022-2026 CBA. However, the compensation structures detailed in the two agreements are not materially different for the purposes of this project. As such, I will refer to both agreements as just “the CBA.”

The CBA breaks player compensation down into three stages: team-mandated, arbitration, and free agency. A player progresses through each stage based on their Major League service (MLS) time. Players receive Major League service time for each day spent on the 26-man roster or the Major League injured list. Each Major League regular season consists of 187 days (typically 183 days prior to 2018) and a player is deemed to have reached "one year" of Major League service upon accruing 172 days in a given year (“Service Time” 2022).

Almost all players with less than two years of service time are in the team-mandated stage. During this stage, teams have “reservation rights” over the player which means they can essentially pay the player whatever they want so long as it is above the minimum salary as determined in the CBA. In 2017, the minimum salary was $535,000 and, due partially to inflation, the minimum salary in 2022 was $700,000 (2017-2021 Collective 2017). Almost invariably, team-mandated players are paid the minimum salary or very close to it. The only source of heterogeneity in these players’ compensation comes from their signing bonuses. Players drafted in the early rounds of the MLB Draft receive significant payouts just for signing a contract with the team who drafted them. For the top draft picks, this payout can be upwards of $5 million, but for the lower draft picks, it can be as little as a few thousand dollars (“Major League Baseball Signing”

2022). During this stage, not only does the player have very little negotiating power over his salary, but his on-field performance has very little impact on his salary as well. The rare exception occurs when a team signs an extremely talented young player to a long-term contract during his first two years because they believe they will have to pay him significantly more if they wait until he reaches the free agency stage ("Service Time" 2022). It should also be noted that the new 2022-2026 CBA established a $50 million bonus pool for players in this stage, but that this pool is not related to a player's contract with his team and is therefore outside of this project's scope (Adams, 2022).

All players with at least three (but less than six) years of Major League service time become eligible for salary arbitration, through which they can earn substantial raises relative to the Major League minimum salary. Designated "Super Two" players, the top 22% - in terms of service time - of players with between two and three years of Major League service time are also eligible for salary arbitration despite having less than three years of service. During the arbitration stage, players and teams negotiate over salaries, primarily based on comparable players who have signed contracts in recent seasons. The team and player exchange salary figures for the upcoming season, and if they can't agree, a panel of arbitrators selects the salary figure of either the player or the club (but not one in between) as the player's salary for the upcoming season. Players eligible for arbitration typically earn significantly higher salaries than team-mandated players because while they are still limited to negotiating only with their team, they can now negotiate for compensation similar to other comparable players in the league. Players remain in the salary arbitration stage until they accumulate six years of Major League service, at which point they become eligible for free agency ("Salary Arbitration" 2022).

Free agency is the final, and most lucrative, stage of MLB player compensation. Free agents are eligible to sign with any club for any terms to which the two parties can agree ("Service Time" 2022). Only at this stage are players able to negotiate contracts with any team in a truly competitive labor market. As a result, not controlling for the different stages of compensation would likely bias the estimates for DRS and OFF compensation effects downward because of the lack of or limited negotiating power held by players in the team-mandated and salary arbitration stages.

Past studies have taken a few different approaches to account for the different contract types. Brown, Link and Rubin (2015) had access to a data source they called "the Joint Exhibit" which is an official summary of contracts compiled by MLB and includes contract information on each player's years of Major League service, length of contract, date of signing, signing bonuses, and base salary. With such detailed information, the authors chose to analyze and report results for players in each of the three contract type groups. However, "the Joint Exhibit" is not publicly available, so other studies have used other approaches. For example, Wasserman (2013) includes an independent variable with the number of MLS years remaining until the player reaches free agency. Clayton and Yermack (2001) and Pollack (2017) includes only those players who have contracts determined through either salary arbitration or free agency because the team-mandated stage does not have any characteristics of a competitive market.

The approach this paper takes is a combination of the Wasserman (2013), Clayton and Yermack (2001), and Pollack (2017) methodologies. I use Major League service time as a proxy for contract stage and create indicator variables for each player, designating them as being in either the team-mandated, arbitration, or free agency groups (Table 1).

Players with MLS less than two years are categorized as team mandated. In addition, to account for "super-two" players, those with between two and three years of MLS and salary less than $50,000 more than the league minimum salary are also categorized as team mandated. Players with MLS greater than or equal to three years and less than six years are put in the arbitration group. In addition, again to account for "super-two" players, those with MLS between two and three years and salary greater than or equal to $50,000 more than the league minimum salary are put in the salary arbitration group. Lastly, players with six or more years of service are free agents.

Table 1
Contract Type Classifications

This table shows the criteria used to classify players as being in either the team mandated, arbitration or free agent contract groups.

Contract Group	Requirements
Team mandated	MLS < 2; OR 2 <= MLS < 3 and salary < league minimum + $50,000
Arbitration	3 <= MLS < 6; OR 2 <= MLS < 3 and salary > league minimum + $50,000
Free agent	MLS >= 6

When testing Hypothesis Two, I drop players in the team-mandated contract stage from the sample. I do this because the highly detailed contract information in "the Joint Exhibit" is not publicly available and I therefore do not have access to signing bonus data for players in the team-mandated contract stage. As such, there is very little heterogeneity in compensation between these players. This decision to exclude players on team mandated contracts is supported by the large differences in salary between each contract type as shown in Table 2. In particular, the standard deviation of salary for players on team mandated contracts is very small when compared to the standard deviations of salary for free agents and arbitration-eligible players. The standard deviation of team mandated salaries is 10.2% of the mean while the standard deviations of free agent and arbitration salaries are 81.9% and 92.0% of their respective means. It is due to this lack of heterogeneity in salary that I choose to exclude players on team mandated contracts when testing my second hypothesis.

Table 2
Salary Descriptive Statistics by Contract Group

This table shows the mean, standard deviation, and median annual salaries for players in the free agent, arbitration, and team mandated contract groups.

Contract Group	Mean	Std. Dev.	Median
Free agent	$ 7,629,534.00	$ 6,250,543.00	$ 5,375,000.00
Arbitration	$ 2,655,986.00	$ 2,442,599.00	$ 1,875,000.00
Team mandated	$ 582,013.50	$ 59,381.91	$ 566,850.00

4. Methodology

A. *Testing Hypothesis One*

To test my first hypothesis that pitch receiving is the most important area of catcher defense, I decompose the DRS statistic into its various components and calculate the standard deviation for each. Consistent with my explanation of "importance" above, the DRS components with the highest standard deviation are the most important to catcher defense because they drive the difference in run saving ability between individual catchers.

B. *Testing Hypothesis Two*

I use two main linear regression models to test my second hypothesis that offensive productivity is compensated more than defensive productivity. Model 1 regresses the holistic performance statistics DRS and OFF against salary. Model 2 breaks DRS and OFF down into their various components and regresses these against salary.

Model 1: $\text{log\_salary}_{ij} = \beta_0 + \text{DRS}_{ij}*\beta_1 + \text{OFF}_{ij}*\beta_2 + \text{FA}_{ij}*\beta_3 + \text{YEAR}_{ij}*\beta_4 + e_{ij}$

Model 2: $\text{log\_salary}_{ij} = \beta_0 + \text{rSZ}_{ij}*\beta_1 + \text{rCERA}_{ij}*\beta_2 + \text{rSB}_{ij}*\beta_3 + \text{rGFP}_{ij}*\beta_4 + \text{rFIELD}_{ij}*\beta_5 + \text{Bat}_{ij}*\beta_6 + \text{UBR}_{ij}*\beta_7 + \text{wGDP}_{ij}*\beta_8 + \text{wSB}_{ij}*\beta_9 + \text{FA}_{ij}*\beta_{10} + \text{YEAR}_j*\beta_{11} + e_{ij}$

The comparability of DRS and OFF makes identifying whether the results support my hypothesis that offensive productivity is over-valued by teams straightforward. Baseball is all about runs. Runs determine the score and consequently winner of every game. Theoretically, creating one run on offense should be just as valuable as saving one run on defense. As such, if the teams are acting rationally and the contract market is efficient, the coefficients for DRS and OFF should be equal. In other words, we would expect teams to pay the same amount for a run gained on offense as a run saved on defense. My hypothesis is that the coefficient for OFF will be higher than the coefficient for DRS (that offense is overvalued, and defense undervalued) because of the misperception that fans enjoy watching offense more than defense. In keeping with this hypothesis, we would expect that the coefficient for Bat will be higher than the coefficients for the components of DRS.

C. *Data*

Data for Defensive Runs Saved (DRS) and OFF (offensive runs above average), the two primary performance statistics of interest, was collected from the online baseball statistics database fangraphs.com. In addition to total DRS, data for each component of DRS was collected as well. Next, contract information including the player’s salary, the number of years of Major League service (MLS) accumulated, team, and player age was sourced from Cot's Baseball Contracts, a Baseball Prospectus database for MLB salary and payroll details.

The sample was constructed by first collecting DRS, DRS components, plate appearances, and OFF statistics for each catcher who had a minimum of at least 90 innings (about 10 games) caught in a given season. This innings qualification was used to ensure that the sample only includes primary catchers (players whose main position is catcher). This dataset was then joined with the salary, MLS, and team data using player name as the key. A total of 59 observations were lost during the join because of contract information being unavailable. This process was repeated for each of the 2016-2022 seasons (Table 3). In total, the dataset has 427 observations and an average of 61 observations per season when including players on team mandated contracts. I will use this sample for testing Hypothesis One. After filtering players with team mandated contracts out, the dataset has 271 observations and is ready for testing Hypothesis Two.

Table 3
Table 3: Dataset Creation

This table shows the filtering and joining process for creating the dataset. OFF and DRS data is sourced from fangraphs.com and salary data comes from Cot's Baseball Contracts.

Year	OFF Obs.	DRS Obs.	Salary Obs.	Joined Obs.	Obs. Lost
2022	89	88	73	61	12
2021	92	83	68	60	8
2020	73	67	69	57	12
2019	91	86	67	63	4
2018	93	89	72	63	9
2017	81	75	69	62	7
2016	88	81	68	61	7
Total	607	569	486	427	59

D. Dependent Variable

The dependent variable in this study is the value of the player's contract. There are several different ways a contract can be valued. The media tends to value contracts by the total amount of money the player will receive over the duration of the contract, or by the average annual value (AAV) which is the total amount divided by the number of years in the contract. Wasserman (2013) bases the valuation of contracts on the present value of each payment from the team to the player, as an average divided by the number of years in the contract. Brown, Link, and Rubin (2015) use the base salary plus prorated signing bonus converted to 2010 dollars using the consumer price index. Pollack (2017) only looks at arbitration contracts (which are one year in length) and free agent contracts signed in the offseason before the season in question. This method avoids looking at salaries from later years in longer contracts that are already locked in which ensures that salaries are a function of relevant market forces. For example, the salary in year four of a contract will not be based upon the player's performance in year three, but rather is still based on their performance several years earlier during the time leading up to the signing of the contract. However, while I believe this approach is reasonable, I also believe that there is a forward-looking nature to contracts. Teams pay players based on how they think they will perform in the future. If a team signs a player to a ten-year contract, they are placing a bet that he will still be playing well in years six and seven. In addition, since I am only studying catchers, there is a limited number of datapoints, so I don't want to constrain the sample to only arbitration and free agent contracts signed in the year before the season in question. As a result, I have decided to use the player's salary at the start of each season as the dependent variable for my models. While salary may not be as detailed a measure as AAV, it allows me to have a much larger sample and I believe it is the best proxy for compensation given the available data.

E. Independent Variables

a. DRS

The two main independent variables of interest are DRS and OFF. For catchers, DRS is simply the sum of its six components described earlier: PART, bunt runs saved, good fielding play runs saved, stolen base runs saved, adjusted earned runs saved, and strike zone runs saved. While each component has its own calculation which can be quite complicated, they all end up being measured in terms of runs saved above or below average. League average is set to 0 runs, so a positive statistic indicates above average performance, and a negative statistic indicates below average. However, due to missing contract data for some players and to filtering out catchers who played less than 90 innings in the season, my dataset does not include every catcher in the league. As a result, I de-mean within my sample so that 0 represents the average for my sample, rather than the league. This doesn't have a huge impact on results because the dataset includes most catchers in the league, but it does slightly raise the average statistic because the catchers excluded from my dataset are those who are less well-known or who got less playing time which indicates they are likely less talented.

Fangraphs.com reports data for total DRS (DRS), good fielding play runs saved (rGFP), stolen base runs saved (rSB), adjusted earned runs saved (cERA), and strike zone runs saved (rSZ). They do not report PART or bunt runs saved, the two components of DRS which essentially make up the "normal fielding" area of catcher defense. However, the four components they do report consistently add up to close, but not quite equal to, total DRS. I assume that this difference is attributable to the omitted PART and bunt runs saved categories. As a result, I created a new variable rFIELD which captures the "normal fielding" area of catcher defense and is equal to the difference in runs saved between total DRS and the sum of the four components that fangraphs.com reports. Both the PART and bunt runs saved calculations follow the same process as the pop-up example given earlier. My new rFIELD category would therefore be calculated using the same methods but can also be derived from the other statistics reported on fangraphs by the following formula: rFIELD = DRS - rGFP - rSB - cERA - rSZ.

The Good fielding plays metric was created by Bill James and defined "a very specific observation of a very narrowly defined event, created in such a way as to keep the scorer's use of judgment to an absolute minimum" (Dewan, "Methodology: Good Play" 2020). For catchers, blocking a pitch thrown in the dirt is the only specific observation that is measured in the GFP statistic. Blocking a pitch is a play that, had it not been made, no one would have faulted the catcher for not making it. It is a pitcher's job to throw the ball near the strike zone and a ball in the dirt is far enough from the strike zone that it is considered an error on the pitcher's part if it gets past the catcher. However, if the catcher does manage to block the ball, he prevents runners from advancing bases which can be converted into runs saved. To calculate a catcher's rGFP, we first find the league average successful block per opportunity for all catchers in a given season by adding up the number of successful blocks across the league and dividing by the total number of block opportunities. We then multiply this league average successful block rate by a specific catcher's number of block opportunities to get his expected number of successful blocks. Next, we find the difference between the expected number of successful blocks to the catcher's actual number of successful blocks. Finally, we multiply the difference by the run value of blocking a pitch in the first (this changes on a yearly basis) and we end with the catcher's rGFP.

Next is rSB or stolen base runs saved which measures a catcher's ability to prevent stolen bases. However, preventing stolen bases is a two-person job. Both the catcher and the pitcher play important roles in the play and the effect of the catcher must be separated from the effect of the pitcher. To do this, *The Fielding Bible* revisits the pitcher's entire history of allowing stolen bases and credits the catcher for every stolen base better (or worse) than the pitcher's career rate. This difference in stolen bases is the number of stolen bases the catcher saved with that specific pitcher on the mound. The calculation is repeated for every pitcher that catcher caught that season and then summed up to get the total number of stolen bases saved. Finally, we use the RE24 matrix to find the run value of a stolen base (around 0.2 runs) and the run value of throwing a runner out (around -0.5 runs). Since each stolen base saved is the difference between a successful stolen base and throwing a runner out, we attribute the run difference (around 0.7 runs) to the catcher. We multiply this runs saved value by the number of stolen bases saved to reach the catcher's rSB statistic (Dewan, "Catcher" 2020).

The Fielding Bible created the statistic adjusted earned runs saved (cERA) to assess a catcher's handling of the pitching staff and pitch calling abilities. The first step in calculating this statistic is to find each pitcher's park-adjusted earned run average (ERA) while throwing to a specific catcher. Next, we take the pitcher's full season park-adjusted ERA, multiply by the number of innings thrown to that particular catcher, and divide by nine to get the expected number of adjusted earned runs allowed by that pitcher/catcher duo. We then subtract the actual number of adjusted earned runs for that tandem to find the number of adjusted earned runs saved for that catcher. Lastly, we multiply this number by a "credibility factor" equal to the total number of innings that catcher caught that season divided three times the number of innings caught in a full season behind the plate. This credibility adjustment is an attempt to address the fact that there is a lot of random noise and variation in why and how a pitcher gives up runs that has nothing to do with the catcher. However, while it is a step in the right direction, there are just so many variables outside of the catcher's control that cERA is still likely somewhat inaccurate (Dewan, "Methodology: Adjusted" 2020).

The last statistic, strike zone runs saved (rSZ), measures what I hypothesize to be the most important area of catcher defense - receiving. The methodology behind this statistic was introduced in a 2015 paper written by Joe Rosales and Scott Spratt that won an award at the Sloan Sports Analytics Conference at MIT in 2015 (Dewan, "Frequently" 2020). Their system Strike Zone Plus/Minus differs from other pitch receiving methodologies in two ways. First, it treats pitchers, batters, and umpires as independent actors in the system rather than treating them as variables to adjust the catcher's performance by. Second, it incorporates data on where the catcher sets his target for the pitch, allowing them to account for the pitcher's command (how close he comes to hitting the target) into the system. To calculate the Strike Zone Plus/Minus for a given pitch, the first step is to calculate the expected strike percentage for that pitch. To do so, pitches are bucketed by pitch location according to a grid that is approximately one inch by one inch. Pitches are then further grouped based on the count they were thrown in, by their horizontal distance from the catcher's target, and by the batter's handedness. With pitches categorized in this way, Rosales and Spratt were able to determine the percent likelihood that each pitch is to be called a strike. If a pitch is called a strike, there is a positive credit to be awarded and if it is called a ball, there is negative credit to be assigned. If a pitch that has a 40% likelihood of being called a strike indeed gets called a strike, then there are 0.6 Strike Zone Plus/Minus points to be allotted to the participants (catcher, pitcher, hitter, umpire) on the pitch. An iterative process involving separating individual tendencies of players and umpires from each other is used to perform this allotment and they end up with the number of "extra strikes" that a catcher earns for his team by his receiving skill. To put this in terms of runs, they use the RE24 matrix to calculate the run expectancy associated with each ball/strike count, and then find the difference in the change in run expectancy between the next pitch being called a ball and the next pitch being called a strike. For the 2010-2013 seasons (the system uses data from the previous four seasons), the average difference in run expectancy between a ball and a strike was .1189 runs. This is multiplied by the number of "extra strikes" the catcher earned to get his rSZ value (Rosales and Spratt 2020).

b. OFF

While there are now several very detailed (sabermetric) statistics used to measure a player's offensive ability such as weighted runs created plus (wRC+), weighted on-base average (wOBA) and weighted runs above average (wRAA), I have chosen to use the offensive runs above average (OFF) statistic due to its comparability to DRS.

In fact, OFF is essentially the offensive counterpart to catcher-specific DRS. It measures a player's context-neutral batting runs and baserunning runs above average and does the best job of crediting a player for the quality and quantity of their total offensive performance during a given period. Like DRS, league average is set to zero, so a positive OFF value denotes above-average performance and a negative OFF denotes below-average performance. It is also a counting statistic (like DRS), so players accrue more (or fewer) runs the more that they play. The formula for OFF is: OFF = batting runs + baserunning runs (Weinberg "Off" 2014).

However, both of these components are composed of several other highly respected statistics. Batting runs is a park-adjusted version of wRAA, which, in its turn, is calculated using wOBA. wOBA combines all the different aspects of hitting into one metric, weighting each aspect in proportion to their actual run value (Slowinski, 2010). It improves upon traditional statistics such as batting average by crediting the hitter for the value of each outcome (single, double, etc.) rather than treating all hits or times on base equally. The formula for wOBA in 2022 was: wOBA = (0.689×uBB + 0.720×HBP + 0.884×1B + 1.261×2B + 1.601×3B +2.072×HR) / (AB + BB – IBB + SF + HBP). Note that the weights on each outcome change on a yearly basis ("WOBA & FIPS Constants" 2022). wRAA is simply wOBA converted into runs and batting runs is simply a park-adjusted version of wRAA. Every ballpark has different dimensions and different geographic locations that affect how easy it is to hit there. For example, closer fences, warmer weather, and higher elevations make hitting home-runs easier. OFF takes these differences into account.

Baserunning runs (BsR) is the sum of the weighted stolen base, ultimate baserunning, and weighted grounded into double play baserunning statistics. Its formula is: BsR = wSB + UBR + wGDP. Each of these three statistics measure a specific aspect of baserunning and combined they quantify the number of runs a player contributes (or costs) their team on offense. When added to batting runs it gets us back to OFF and provides us with a comprehensive measure of a player's offensive output.

c. Contract Type

Next, I include control variables for contract type. ARB and FA are both dummy variables that indicate if the player is in the salary arbitration stage or free agency stage respectively. These variables are created using Major League service time information as described above. In the specification with Team-mandated contract players included, I will include both the ARB and FA dummy variables and let the Team-mandated category be the excluded dummy variable which the ARB and FA coefficients will be relative to. For the specification without team-mandated contract players, I will let ARB be the excluded dummy variable that the FA coefficient will be relative to.

d. Year Effects

The last independent variable included is a set of year fixed effects. This is to control for inflation as well as the incremental increases in minimum salary stipulated by the CBA.

5. Exploratory Data Analysis

Table ***4*** and 5 below report the descriptive statistics for the two primary datasets used in this study. The first includes players on team mandated contracts and is used for testing Hypothesis One and the second excludes players on team mandated contracts and is used to test Hypothesis Two.

Table 4
Descriptive statistics including players on team mandated contracts (427 obs.)
This table shows the mean, standard deviation, and median for salary and log salary, MLS and the contract type dummies, innings and plate appearances, DRS and its components, and OFF and its components in the dataset that includes team mandated players used for testing Hypothesis One.

Variable	Mean	Std. Dev.	Median
MLS	4.28	3.5	3.16
Salary	3412477	4726838	1500000
log_salary	14.36	1.11	14.22
Inn	538.51	283.83	504.1
DRS	0	7.01	-0.37
rSZ	0	4.65	-0.2
rCERA	0	2.74	0
rSB	0	2.26	-0.04
rGFP	0	2.29	-0.12
rFIELD	0	1.34	-0.01
OFF	0	8.84	-0.39
Bat	0	8.68	-0.71
UBR	0	1.47	0.27
wGDP	0	0.82	0.1
wSB	0	0.39	0.05
FA	0.3	0.46	0
ARB	0.33	0.47	0
MAND	0.37	0.48	0

As discussed above, the performance statistics are standardized to have mean equal to zero so that positive values can be interpreted as above average performance and negative values as below average. As a result, standard deviation is the most meaningful measure for determining the importance of the performance statistics. I will address the standard deviations of the components of DRS later during discussion of the Hypothesis One results but will examine DRS and OFF here. DRS has a standard deviation of 7.01 runs and OFF has a standard deviation of 8.84 runs in this sample. This means that there is more variation in offensive ability than there is in defensive ability. This is important because it means that an elite hitting catcher contributes more runs to a team over an average hitting catcher than an elite defensive catcher contributes over an average defensive catcher. There is essentially a wider skill gap between catchers on offense than there is on defense. Based on this finding, one might expect that teams would be willing to pay elite offensive catchers more than they would be willing to pay elite defensive catchers. In particular, examining the standard deviations of the components of OFF, we see that batting is by far the greatest source of variation. In addition, the table shows that the average player in the dataset has been in the league for 4.28 years, catches 538.1 innings in a season, and has a yearly salary of $3,412,477. Finally, we can see that 30% of the players were on a free agent contract, 33% were on an arbitration contract, and 37% were on a team-mandated contract.

Table 5
Descriptive Statistics excluding players on team mandated contracts (270 obs.)

This table shows the mean, standard deviation, and median for salary and log salary, MLS and the contract type dummies, innings and plate appearances, DRS and its components, and OFF and its components in the dataset that excludes team mandated players used for testing Hypothesis Two.

Variable	Mean	Std. Dev.	Median
MLS	6.16	3.05	5.15
Salary	$5,041,821	$5,287,418	$3,000,000
log_salary	14.99	0.93	14.91
Inn	573.93	286.24	559.1
DRS	0.06	7.12	-0.37
rSZ	0.05	4.81	-0.2
rCERA	0.01	2.73	0
rSB	-0.09	2.38	-0.04
rGFP	0.16	2.48	-0.12
rFIELD	-0.07	1.3	-0.01
OFF	-0.34	9.26	-0.69
Bat	-0.12	9.01	-0.81
UBR	-0.13	1.55	0.07
wGDP	-0.08	0.89	0
wSB	-0.02	0.43	0.05
FA	0.48	0.5	0
ARB	0.52	0.5	1

It is interesting to compare Table 5 with Table 4 as it illustrates the differences between players on free agent and arbitration contracts and players on team mandated contracts. The primary difference between the two datasets is in mean salary. When team mandated contracts are included, the mean salary is $3,412,477. When they are excluded, the mean salary jumps over $1.5 million to $5,041,821. One would expect that the more well-established, experienced, and higher paid players on free agent and arbitration contracts would, on average, have above average performance statistics. However, comparing means in Table 4 and Table 5 shows that the more experienced, higher-paid players in fact have slightly below-average performance in several areas, most notably in every offensive statistic. However, while older players appear to be slightly worse hitters, they also appear to be slightly better defenders as seen by the positive means for DRS, rSZ (pitch receiving), and rGFP (blocking). Since we know that teams pay free agents and arbitration-eligible players more, this quick comparison of descriptive statistics hints that teams might value defensive productivity more than offensive productivity at the catcher position. However, let's examine the regressions results from the main specifications for Hypothesis Two before drawing any conclusions.

Table 6
Correlation Matrix

This table shows the correlation matrix (Pearson coefficients) of MLS, Salary, DRS and its components and OFF and its components. These coefficients are calculated on a univariate basis not controlling for other factors.

Variable	MLS	Salary	DRS	rSZ	rCERA	rSB	rGFP	rFIELD	Off	Bat	UBR	wGDP	wSB
MLS	1												
Salary	0.605	1											
DRS	-0.091	0.041	1										
rSZ	-0.066	0.071	0.731	1									
rCERA	-0.044	-0.074	0.487	0.123	1								
rSB	-0.079	-0.008	0.4	0.001	0.021	1							
rGFP	0.032	0.099	0.47	0.091	0.042	0.137	1						
rFIELD	-0.08	-0.037	0.224	-0.053	-0.026	0.127	0.117	1					
Off	-0.042	0.169	-0.021	0.012	-0.085	0.041	-0.043	0.028	1				
Bat	0.005	0.223	-0.02	0.014	-0.093	0.036	-0.033	0.032	0.975	1			
UBR	-0.165	-0.18	-0.049	-0.039	-0.018	0.01	-0.04	-0.036	0.14	-0.058	1		
wGDP	-0.204	-0.225	0.044	0.03	0.089	-0.007	-0.06	0.056	0.181	0.052	0.195	1	
wSB	-0.008	0.034	0.073	0.058	0.015	0.099	0.025	-0.06	0.066	-0.033	0.262	0.108	1

Examining the correlation matrix given in Table 6, there are a few key relationships that jump out. First, as a quick check, MLS and salary have a .605 correlation coefficient which aligns with our understanding and interpretation of MLB contract types - as players accrue years of service, they become eligible for more lucrative contracts. However, despite the strong correlation between MLS and salary, there is very little correlation between MLS and Off and MLS and DRS. This suggests that more veteran catchers aren't performing any better on the field than their less-experienced counterparts.

The next correlation coefficients of note are those between salary and DRS and between salary and OFF. The correlation between salary and DRS is 0.041 which is not statistically different from 0 (p-value = 0.3978) suggests that a catcher's compensation is not closely tied to his defensive performance. This is surprising as defense is not only an important part of the game of baseball but is also particularly important to the demanding and highly specialized catcher position. As for the correlation between salary and OFF, the coefficient of 0.169 is statistically different from 0 (p-value = 0.0004), but still does not suggest an especially strong relationship. While the fact that the correlation between salary and OFF is higher aligns with my hypothesis that offensive productivity is more highly compensated than defensive productivity, it is still surprising that the relationship is not stronger. One reason for the weak correlation for both OFF and DRS and salary is that there isn't enough variation in skill between catchers to justify teams spending significantly more on one catcher versus another.

Additional key relationships to point out are those between the holistic performance statistics (DRS and OFF) and their most highly correlated component. DRS is most highly correlated with rSZ, the statistic measuring receiving. They have a strong correlation coefficient of 0.731 which is .244 greater than the next most correlated component. While this does not directly support my hypothesis that rSZ is the source of greatest variation in catcher defense, it is reassuring to see that it is also the most highly correlated with DRS. As for OFF, it is most highly correlated with the Bat component which measures a catcher's runs created while batting. At 0.975, this relationship is extremely strong, but not particularly surprising. The other three components of OFF measure baserunning ability which plays a much smaller role in creating runs as baserunning opportunities depend on players being able to get on base in the first place. A player could be the fastest person in the world, but if they can't hit the ball or get on base, it won't help their team score runs.

A final relationship to explore is the correlation, or lack thereof, between DRS and Off. One might reason that these two variables might be negatively correlated, as offensive prowess might make up for a lack of defensive ability and vice versa. However, these two variables have a correlation coefficient of -0.021. While the sign is in the right direction, the relationship is extremely weak and suggests that better offensive catchers do not perform significantly worse on defense and better defensive catchers do not perform significantly worse on offense.

6. Results

A. Hypothesis One Results:

To test the hypothesis that receiving accounts for the greatest difference in catchers' defensive performances, I decomposed DRS into its various components and calculated each component's standard deviation. These are reported in Table 7 below.

Table 7
Standard deviation of DRS and its components

This table shows the standard deviation for DRS and each of its components. These values are calculated using the dataset that includes team mandated players.

Variable	Std. Dev.
DRS	7.01
rSZ	4.65
rCERA	2.74
rGFP	2.29
rSB	2.26
rFIELD	1.34

From Table 7 we can see that, at 4.65 runs saved, rSZ (pitch receiving) has the highest standard deviation of all the DRS components and, using the F-test, is statistically significantly higher than the next component. rSZ is followed by rCERA (handling of the pitching staff and pitch-calling) at 2.74 runs saved, rGFP (blocking) at 2.29 runs saved, rSB (preventing stolen bases) at 2.26 runs saved, and lastly rFIELD (fielding bunts, catching pop-ups, and tagging runners out at the plate) at 1.34 runs saved.

For my first hypothesis, I proposed that pitch receiving, as measured by the strike zone runs saved (rSZ) component of DRS, is the greatest source of variation in defensive performance for MLB catchers. This hypothesis is supported by the results in Table 7. The standard deviation of strike zone runs saved (rSZ) is nearly 70% greater than the standard deviation of the next closest component rCERA. This means that the difference between elite and average catchers in receiving is greater than the difference between elite versus average catchers in the other components. These results imply that teams should prioritize acquiring and compensating more highly catchers who are elite receivers because over the course of the season this ability will save the team more runs than if they signed a catcher who is elite in another defensive component. This also suggests that younger catchers should focus their practice time on receiving, as becoming a better receiver is likely the best path to distinguish oneself from other catchers. These findings are consistent with recent literature on catcher defense and the importance of receiving. With the introduction of pitch-tracking technology such as PitchF/X and Statcast, studies done by Fast (2011), Rosales and Spratt (2015), and Hook (2020) were able to successfully quantify the pitch receiving skill and found that this ability could earn or lose a team an additional two wins per season.

B. Hypothesis Two Results:

Table 8
Model 1 Results

This table shows the regression results from Model 1. DRS and OFF are the primary variables of interest and FA and the year variables are important controls.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5607	0.1314	110.816	0.0000 ***
DRS	0.0169	0.0068	2.4698	0.0142 **
OFF	0.0175	0.0052	3.3651	0.0009 ***
FA	1.0225	0.0975	10.4922	0.0000 ***
year_2022	0.0725	0.1759	0.4121	0.6806
year_2021	-0.1446	0.1786	-0.8096	0.4189
year_2020	-0.0394	0.1807	-0.2179	0.8277
year_2019	-0.1485	0.1751	-0.8485	0.3969
year_2018	0.0768	0.1774	0.4328	0.6655
year_2017	-0.2167	0.1727	-1.2553	0.2105

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

Residual standard error: 0.7761 on 261 degrees of freedom
Multiple R-squared: 0.3314, Adjusted R-squared: 0.3084
F-statistic: 14.38 on 9 and 261 DF, p-value: < 2.2e-16

Table 8 reports the coefficient estimate, standard error, and statistical significance for each variable in Model 1. As can be seen in the table, free agents earn a 102.25% greater salary on average than arbitration eligible players and that is statistically significant at the 1% level. Next, a one run increase of OFF implies a 1.75% increase in salary on average that is significant at the 1% significance level. Similarly, a one run increase in DRS implies a 1.69% increase in salary on average that is significant at the 5% level. These estimates are not statistically significantly different (Wald test statistic = -.06).

My second hypothesis was that teams value offensive productivity more than they value defensive productivity because of the perception that offense is more fun to watch, driving up fan attendance and total revenue. Based on the regression results in Table 8 and Table 9, this hypothesis is rejected. Since the DRS and OFF statistics (and their components) are measured in terms of runs above average and are standardized to have a mean of 0, their coefficients can be directly compared and interpreted as the percent change in salary caused by a one run increase in the statistic. As such, the results show that teams in fact increase salary by almost the same percentage for a one run increase in DRS as they do for a one run increase in OFF. A one run increase in DRS causes a 1.69% increase in salary while a one run increase in OFF causes a 1.75% increase in salary. This shows that teams are indeed efficient with their compensation of catchers from an on-field performance standpoint.

Table 9
Model 2 Results

This table shows the regression results from Model 2. The components of DRS and the components of OFF are the primary variables of interest and FA and the year variables are important controls.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5631	0.1263	115.328	0.0000 ***
rSZ	0.029	0.0098	2.9597	0.0034 ***
rCERA	-0.0171	0.0171	-1.0004	0.3181
rSB	0.0179	0.0197	0.909	0.3642
rGFP	0.0227	0.0187	1.2151	0.2255
rFIELD	0.0023	0.0359	0.065	0.9482
Bat	0.0212	0.0052	4.0707	0.0001 ***
UBR	-0.0621	0.0321	-1.9384	0.0537 *
wGDP	-0.1771	0.0532	-3.3282	0.0010 ***
wSB	0.0496	0.112	0.4433	0.6579
FA	0.9418	0.095	9.9153	0.0000 ***
year_2022	0.0656	0.1686	0.3892	0.6975
year_2021	-0.1321	0.1733	-0.7619	0.4468
year_2020	0.0406	0.1735	0.2341	0.8151
year_2019	-0.1476	0.1684	-0.8766	0.3815
year_2018	0.0485	0.1717	0.2824	0.7779
year_2017	-0.2029	0.1655	-1.2258	0.2214

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1

Residual standard error: 0.7398 on 254 degrees of freedom
Multiple R-squared: 0.4087, Adjusted R-squared: 0.3714
F-statistic: 10.97 on 16 and 254 DF, p-value: < 2.2e-16

Table 9 reports the coefficient estimate, standard error, and statistical significance for each variable in Model 2. As can be seen in the table, free agency is still statistically significant at the 1% level and, on average, increases salary by 94.18% over arbitration eligible players. In addition, a one run increase in rSZ implies a 2.90% increase in salary on average and is significant at the 1% level. On the offensive performance side, a one run increase in Bat implies a 2.12% increase in salary on average and is statistically significant at the 1% level. Also on the offensive side, a one run increase in wGDP (grounding into double plays) implies a -17.71% change in salary on average and this is statistically significant at the 1% level. Lastly, a one run increase in UBR (ultimate baserunning) implies a -6.21% change in salary on average and this is significant at the 10% level. The negative effects on salary associated with the baserunning components wGDP and UBR imply that teams do not value productivity in these areas for catchers. While of course teams would prefer a fast catcher over a slow catcher all else being equal, there is likely a complicated correlation between these baserunning components and other DRS and OFF-related components which the model has trouble parsing out. What is likely occurring is that the most highly compensated catchers are more prone to grounding into double plays and are poor baserunners. This makes sense as good sprint speed and acceleration are not integral to being an elite catcher - other skills such as hand-eye coordination and flexibility are far more important. As a result, the model reports negative coefficients for these two variables even though in reality a team would likely pay more for a catcher with higher wGDP and UBR, all else being equal.

The results from Model 2 show that teams pay catchers a 2.90% higher salary on average for each run they save by pitch receiving. This is the most compensated component of DRS, which is expected given the result from Hypothesis One that rSZ is the greatest source of variation between catchers on defense. In addition, the 2.90% increase in salary from an additional one run of rSZ is greater than the 2.12% salary increase from an additional one run of batting. Both the coefficients on rSZ and Bat are statistically significant and suggest that teams are in fact compensating catchers slightly more for their pitch receiving productivity than their batting productivity. This finding means that teams do not inaccurately compensate catchers based on the misperception that offensive productivity is more valuable because it is more exciting for fans. The results indicate that, if anything, teams slightly over-compensate catchers for their pitch receiving abilities and under-compensate them for their hitting. This differs from prior literature such as Maurice (2010) and Ehrlich and Potter (2020) that found that teams pay more for offensive performance than they do for defensive performance. However, these prior studies look at all positions, rather than at the catcher position specifically. This is an important difference because of the specialized nature and importance of the catcher position on defense. Other than perhaps the pitcher, the catcher has the greatest impact on the game when a team is on defense. As such, it is reasonable that teams prioritize defensive productivity at the catcher position more than at other positions when making roster and salary decisions.

This finding is not entirely surprising as every team in the MLB now has a data analytics department whose entire purpose is to accurately value players (Avidon, 2022). With such a significant investment in evaluating and projecting player performance, it would be surprising if the market were inefficient given the data currently available today. The only inefficiencies that remain are likely in the areas of catching defense that are not yet adequately measured, such as pitch-calling and leadership.

7. Conclusion

This paper's two main goals are to determine which area of catcher defense is the greatest differentiator between catchers' defensive productivity and to determine how teams are compensating catchers for their defensive and offensive productivity. With these goals in mind, the driving hypotheses of the analysis are that pitch receiving is the greatest source of variation in defensive productivity for catchers and that offensive productivity is compensated more than defensive productivity. The first hypothesis is supported by the results in Table 7. The standard deviation of strike zone runs saved is 4.65 runs, nearly 70% greater than the standard deviation of the next closest component rCERA.These results imply that teams should prioritize signing, and should be willing to pay more for, catchers who are elite receivers because this ability will save the team more runs over the course of the season than if they sign a catcher who is elite in another defensive component.

The second hypothesis is rejected based on the regression results in Table 8 and Table 9. Table 8 shows that teams in fact increase salary by almost the same percentage for a one run increase in DRS as they do for a one run increase in OFF. In particular, a one run increase in DRS causes a 1.69% increase in salary while a one run increase in OFF causes a 1.75% increase in salary. Breaking down DRS and OFF into their respective components yielded the results that a one run increase in rSZ implies a 2.9% increase in salary on average and a one run increase in Bat implies a 2.12% increase in salary on average. These results show that teams are in fact compensating catchers slightly more for their pitch receiving productivity than their batting productivity. This finding means that teams do not inaccurately compensate catchers based on the misperception that offensive productivity is more valuable because it is more exciting for fans. In fact, this paper finds that teams are quite efficient with their compensation of catchers from an on-field performance standpoint. This differs from the findings of prior literature that, in analyzing all field positions, concluded that offensive productivity is compensated more than defensive productivity. This implies that teams prioritize defensive productivity at the catcher position more than at other positions when making roster and salary decisions. Overall, this study finds that MLB teams compensate catchers efficiently in terms of their on-field performance. Given that sabermetrics is now widely accepted, it makes sense that teams are efficient since they all have access to this same data. However, it should be noted that all aspects of the game are not yet easily measurable. Pitch-calling and leadership are two skills that are important to have in an elite catcher that are still difficult to accurately quantify. As such, there is still currently a limit to what sabermetrics can do. There is yet room for inefficiencies in compensation to occur, and room for scouts and other front-office managers to identify players who might be under-valued.

A. Limitations

It should be noted that this study has its limitations. Chief among them is the unavailability of detailed compensation information. While salary is an adequate proxy for total compensation, the inclusion of total contract value, contract length, signing bonus, and other incentives information would make for a much richer dependent variable. For example, a contract paying $10 million a year for 10 years is much more valuable than a single-year contract that pays $10 million because it guarantees the player long-term financial security. Signing bonus and other incentives are also important factors when measuring a player's total compensation. For young, recently drafted players especially, signing bonus can make up the entire difference in compensation between these unproven players. Armed with this information, researchers would be able to include players on team mandated contracts which would make the dataset more robust by increasing the number of observations.

Without this additional detail about a player's total compensation, the results of this study are likely less accurate than the results of a study using a more holistic measure of compensation. This limitation likely dulls the estimated effects because incentives are often tied to meeting certain performance benchmarks. However, it is hard to say how this limitation would affect the difference in estimates for DRS, OFF and their respective components because incentives can be tied to performance on both the defensive and offensive sides of the game.

B. *Extensions*

Possible extensions to this study include adding more years of data and extending the analysis to positions other than catcher. Including more years of data could prove fruitful as it would increase the number of observations in the dataset. One could also try different spans of time to see if there was ever a period, potentially in the early-mid

2010’s, when pitch-tracking technology had not yet been widely accepted, where catcher compensation in the MLB wasn’t efficient.

A second interesting extension to this study would be to repeat the analysis for other field positions. For example, it would be interesting to see if teams still prioritize defensive productivity at less-specialized positions like outfield and first and second base. I would hypothesize that the variation in defensive productivity at the corner outfield positions is much less than the variation in defensive productivity at the catcher position, and as a result offensive productivity for outfielders would be compensated more highly than defensive productivity.

8. Appendix

A. Robustness Check: Effects of Covid-19

Like other professional sporting leagues, Major League Baseball was significantly affected by the Covid19 pandemic. The 2020 MLB season was shortened from its usual 162-game regular season to only a 60-game regular season which ran from July 23rd through October 28th (Feinsand, 2020). This shortened season means that catchers in 2020 had fewer games to differentiate themselves from each other in terms of the DRS and OFF statistics. Better catchers will, on average, perform better each game than worse catchers. Over the course of the season, the runs saved and runs created that they contribute to or lose their team accrue and there is a wider range of values. Since there were significantly fewer games played in 2020, this range is smaller than the other seasons and this can be seen by comparing the standard deviation of the sample including the 2020 season and the standard deviation of the sample excluding the 2020 season.

App. Table 1
Comparing Standard Deviations of DRS and OFF

This table shows the Standard Deviations of DRS and OFF in 2020 and all years in the period of analysis except for 2020.

	2020	All years except for 2020
DRS Std. Dev.	3.44	7.41
OFF Std. Dev.	4.96	9.19

The standard deviation of DRS in 2020 is less than half of the average standard deviation of DRS over the other six seasons in the sample. The standard deviation of OFF is also much smaller than the average OFF in the other seasons. As a result, including the 2020 season lowers both statistics’ standard deviations by about 0.4 runs.

The most likely area this could cause issues in is Hypothesis One. As described above, the components of DRS measure different areas of catching defense. Catchers save runs in these areas in different ways. Runs saved by receiving pitches well (rSZ) are accrued incrementally pitch-by-pitch. In contrast, runs saved by preventing stolen bases (rSB) are accumulated via the comparatively rare occurrence of runners attempting to steal a base. When these plays do happen however, the run impact is much higher than getting a single pitch called a strike instead of a ball. So, since 2020 has a much smaller sample size of games, there is a chance that during those games there were more (or fewer) stolen base attempts than an average 162 game season and this could cause a disproportionate number of runs saved to be attributed to the rSB component. As such, to test that my results for Hypothesis One are robust to the inclusion of the 2020 season, I report the standard deviations of the DRS components with and without the 2020 season.

App. Table 2
Comparing Standard Deviations of DRS and DRS Components

This table shows the Standard Deviations of DRS and the DRS components when including versus excluding 2020 observations.

Component	Std. dev. with 2020	Std. dev. without 2020
rSZ	4.65	4.95
rGFP	2.29	2.42
rCERA	2.74	2.86
rSB	2.26	2.38
rFIELD	1.34	1.41
DRS	7.01	7.41

The standard deviations of the components of DRS are not changed by substantially different amounts with the inclusion of 2020 data. All components increase by around 5% when dropping observations from 2020. In addition, rSZ still has the largest standard deviation by nearly double the next component. From this it can be concluded that the Hypothesis One results are robust to the effects of the shortened 2020 season caused by Covid-19.

As for Hypothesis Two results, the effects of Covid-19 are minimal because there are just not that many new contracts signed year-to-year, and any differences in salaries between 2020 and the other years is picked up by the year fixed effects. The results for Model 1 and Model 2 excluding 2020 observations are tabulated below.

App. Table 3
Model 1 Results Excluding 2020 Observations

This table shows the regression results from Model 1 when excluding observations from the 2020 season.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5415	0.1331	109.2371	0.0000 ***
DRS	0.0172	0.0070	2.4625	0.0145 **
Off	0.0175	0.0053	3.2693	0.0012 ***
FA	1.0483	0.1055	9.9411	0.0000 ***
year_2022	0.0732	0.1769	0.4140	0.6793
year_2021	-0.1467	0.1797	-0.8167	0.4150
year_2019	-0.1506	0.1761	-0.8552	0.3934
year_2018	0.0738	0.1785	0.4136	0.6796
year_2017	-0.2178	0.1737	-1.2542	0.2111

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1
Residual standard error: 0.6227 on 416 degrees of freedom
Multiple R-squared: 0.6948, Adjusted R-squared: 0.6875
F-statistic: 94.71 on 10 and 416 DF, p-value: < 2.2e-16

App. Table 4
Model 2 Results Excluding 2020 Observations

This table shows the regression results from Model 2 when excluding observations from the 2020 season.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5502	0.1278	113.8297	0.0000 ***
rSZ	0.0260	0.0100	2.6118	0.0096 ***
rCERA	-0.0109	0.0176	-0.6164	0.5383
rSB	0.0077	0.0204	0.3789	0.7051
rGFP	0.0255	0.0191	1.3316	0.1844
rFIELD	0.0161	0.0371	0.4346	0.6643
Bat	0.0216	0.0054	4.0313	0.0001 ***
UBR	-0.0662	0.0332	-1.9936	0.0474 **
wGDP	-0.1811	0.0546	-3.3175	0.0011 ***
wSB	0.1053	0.1156	0.9111	0.3632
FA	0.9591	0.1024	9.3687	0.0000 ***
year_2022	0.0711	0.1692	0.4204	0.6746
year_2021	-0.1318	0.1741	-0.7572	0.4498
year_2019	-0.1388	0.1691	-0.8212	0.4124
year_2018	0.0633	0.1725	0.3668	0.7141
year_2017	-0.1956	0.1661	-1.1770	0.2405

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

Residual standard error: 0.6012 on 409 degrees of freedom
Multiple R-squared: 0.7204, Adjusted R-squared: 0.7087
F-statistic: 61.98 on 17 and 409 DF, p-value: < 2.2e-16

From the above tables we can see that results do not change significantly when 2020 observations are excluded. While the magnitude of the coefficients decrease, their relationships to each other and their statistical significance do not. This implies that the distribution of the composite DRS and OFF variables don't meaningfully change when excluding 2020.

B. Robustness Check: Including Team Mandated Players

In addition to checking that the results are robust to the effects of Covid-19, I also check that they are robust to the inclusion of players on team mandated contracts. As described earlier, I chose not to include players on team mandated contracts in the dataset for testing Hypothesis Two. This decision was motivated by the lack of heterogeneity in salary for players on team mandated contracts. This lack of heterogeneity is a sign that these players' contracts do not change significantly based upon their on-field performance and are based more on the fact that they are not yet eligible for arbitration. However, I believe it is still pertinent to check that including these players doesn't change results significantly. Below, I report the results from Model 1 and Model 2 when run on the dataset including team mandated players.

App. Table 5
Model 1 Results Including Team Mandated Players

This table shows the regression results for Model 1 when including observations designated as being team mandated players.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	13.2728	0.0893	148.6024	0.0000 ***
DRS	0.0113	0.0044	2.5895	0.0099 ***
Off	0.0130	0.0035	3.7335	0.0002 ***
FA	2.2564	0.0743	30.3567	0.0000 ***
ARB	1.2480	0.0731	17.0836	0.0000 ***
year_2022	0.1199	0.1130	1.0613	0.2892
year_2021	-0.0545	0.1135	-0.4798	0.6316
year_2020	-0.0003	0.1159	-0.0028	0.9978
year_2019	-0.0610	0.1119	-0.5451	0.5859
year_2018	0.0544	0.1120	0.4854	0.6276
year_2017	-0.1428	0.1125	-1.2695	0.2050

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

Residual standard error: 0.6227 on 416 degrees of freedom
Multiple R-squared: 0.6948, Adjusted R-squared: 0.6875
F-statistic: 94.71 on 10 and 416 DF, p-value: < 2.2e-16

In this model, we can see that being an arbitration eligible player implies a 124.8% higher salary on average than a team mandated player and that being a free agent implies a 225.64% higher salary than a team mandated player. In addition, the R-squared value jumped from around 33% up to nearly 70%. This comes as a result of the additional 157 observations in the dataset. However, as can be seen, these results are very similar to the main results which exclude team mandated players. The estimates for DRS and OFF are still quite similar to each other and still statistically significant which implies that teams are still generally efficient with their compensation of players based on their on-field performance.

App. Table 6
Model 2 Results Including Team Mandated Players

This table shows the regression results for Model 1 when including observations designated as being team mandated players.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	13.3072	0.0867	153.4374	0.0000 ***
rSZ	0.0192	0.0065	2.9748	0.0031 ***
rCERA	-0.0124	0.0109	-1.1393	0.2552
rSB	0.0097	0.0133	0.7319	0.4646
rGFP	0.0152	0.0131	1.1612	0.2462
rFIELD	0.0089	0.0224	0.3968	0.6918
Bat	0.0159	0.0034	4.6198	0.0000 ***
UBR	-0.0464	0.0214	-2.1715	0.0305 **
wGDP	-0.1342	0.0371	-3.6137	0.0003 ***
wSB	0.0312	0.0783	0.3990	0.6901
FA	2.1703	0.0740	29.3457	0.0000 ***
ARB	1.2233	0.0712	17.1936	0.0000 ***
year_2022	0.1138	0.1094	1.0400	0.2990
year_2021	-0.0538	0.1103	-0.4878	0.6259
year_2020	0.0481	0.1124	0.4282	0.6687
year_2019	-0.0742	0.1082	-0.6860	0.4931
year_2018	0.0348	0.1085	0.3207	0.7486
year_2017	-0.1490	0.1089	-1.3688	0.1718

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

Residual standard error: 0.6012 on 409 degrees of freedom
Multiple R-squared: 0.7204, Adjusted R-squared: 0.7087
F-statistic: 61.98 on 17 and 409 DF, p-value: < 2.2e-16

The results from Model 2 are also very similar to the main results. A one run increase in rSZ is still compensated slightly more than a one run increase in Bat. As a result, I feel confident that the main results are robust to the inclusion of team mandated players.

C. Alternate Specifications Results

Model 3: log_salary = β_0 + rSZ*β_1 + rCERA* β_2+ rSB*β_3 + rGFP*β_4 + rFIELD*β_5 + OFF*β_6 + FA*β_7 + YEAR*β_8 + e

App. Table 7
Model 3 Results

This table shows the regression results for Model 3, an alternative specification to my main tests for Hypotheses One and Two. Model 3 regresses OFF, the components of DRS, FA and year effects on log(salary). The results are similar to those from Models 1 and 2. Estimates for rSZ and OFF are statistically significant and the rSZ effect is larger.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5549	0.1315	110.7154	0.0000 ***
rSZ	0.0315	0.0101	3.1148	0.0020 ***
rCERA	-0.0247	0.0178	-1.3865	0.1668
rSB	0.0156	0.0204	0.7672	0.4437
rGFP	0.0257	0.0195	1.3158	0.1894
rFIELD	0.0060	0.0371	0.1614	0.8719
OFF	0.0154	0.0053	2.9122	0.0039 ***
FA	1.0250	0.0971	10.5520	0.0000 ***
year_2022	0.0771	0.1758	0.4389	0.6611
year_2021	-0.1171	0.1783	-0.6570	0.5117
year_2020	-0.0257	0.1802	-0.1427	0.8867
year_2019	-0.1464	0.1753	-0.8355	0.4042
year_2018	0.0590	0.1776	0.3323	0.7399
year_2017	-0.2331	0.1719	-1.3559	0.1763

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1

Residual standard error: 0.7715 on 257 degrees of freedom
Multiple R-squared: 0.3493, Adjusted R-squared: 0.3164
F-statistic: 10.61 on 13 and 257 DF, p-value: < 2.2e-16

Model 4: log_salary = β_0 + DRS*β_1 + Bat*β_2 + UBR*β_3 + wGDP*β_4 + wSB*β_5 + FA*β_6 + YEAR*β_7 + e

App. Table 8
Model 4 Results

This table shows the regression results for Model 4, an alternative specification to my main tests for Hypotheses One and Two. Model 4 regresses DRS, the components of OFF, FA and year effects on log(salary). The results are like those from Models 1 and 2. Estimates for DRS, Bat, UBR, wGDP are statistically significant.

Variable	Estimate	Std.Error	t-Statistic	p.Value
(Intercept)	14.5664	0.1258	115.7607	0.0000 ***
DRS	0.0166	0.0066	2.5364	0.0118 **
Bat	0.0231	0.0051	4.5491	0.0000 ***
UBR	-0.0679	0.0318	-2.1357	0.0336 **
wGDP	-0.1786	0.0532	-3.3583	0.0009 ***
wSB	0.0625	0.1115	0.5605	0.5757
FA	0.9355	0.0949	9.8541	0.0000 ***
year_2022	0.0631	0.1681	0.3752	0.7078
year_2021	-0.1511	0.1731	-0.8725	0.3838
year_2020	0.0352	0.1735	0.2028	0.8395
year_2019	-0.1458	0.1676	-0.8696	0.3853
year_2018	0.0640	0.1707	0.3750	0.7080
year_2017	-0.1866	0.1656	-1.1271	0.2607

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.

Residual standard error: 0.7415 on 258 degrees of freedom
Multiple R-squared: 0.3966, Adjusted R-squared: 0.3685
F-statistic: 14.13 on 12 and 258 DF, p-value: < 2.2e-16

D. Individual Catcher Rankings

App. Table 9

Catcher Rankings by Player

This table lists the all the catchers in the sample including team mandated players. Avg. total performance is simply the sum of the catcher's average DRS and average OFF statistics. Average salary is the average of the catcher's salary over the period of analysis (2016-2022). Likewise, seasons played is the number of seasons that catcher played more than 90 innings between 2016-2022.

Player Name	Avg. Tot. Perf.	Avg. DRS	Avg. OFF	Avg. Salary	Seasons Played
Cal Raleigh	22.6	14.0	8.6	$ 702,900	1
Will Smith	17.3	3.3	14.0	$ 631,333	3
Buster Posey	15.1	8.8	6.3	$ 21,897,778	5
J.T. Realmuto	14.2	1.0	13.2	$ 9,110,000	7
David Ross	13.0	12.0	1.0	$ 2,500,000	1
Sean Murphy	10.5	4.3	6.1	$ 621,333	3
Yasmani Grandal	10.5	6.4	4.0	$ 12,421,429	7
William Contreras	10.0	-4.0	14.0	$ 710,000	1
Alejandro Kirk	8.0	3.0	5.0	$ 641,950	2
Willson Contreras	7.7	1.7	6.0	$ 3,769,667	6
Tyler Stephenson	5.3	0.0	5.3	$ 645,250	2
Tyler Flowers	4.9	5.4	-0.5	$ 3,400,000	5
Manny Pina	4.2	7.6	-3.4	$ 1,239,660	5
Austin Barnes	3.2	5.3	-2.1	$ 1,178,333	6
Mitch Garver	3.2	-3.0	6.2	$ 1,390,500	5
Garrett Stubbs	2.5	-3.0	5.5	$ 704,500	1
Mike Zunino	2.1	4.3	-2.3	$ 3,576,250	6
Danny Jansen	1.5	3.3	-1.7	$ 922,850	4
Roberto Perez	1.4	9.4	-8.1	$ 2,830,829	7
Jose Trevino	1.2	9.7	-8.5	$ 623,167	3
Russell Martin	0.9	5.0	-4.1	$ 18,750,000	4
Salvador Perez	0.6	-0.5	1.1	$ 10,216,667	6
Alex Avila	0.5	2.2	-1.6	$ 3,083,333	6
Jacob Stallings	0.5	8.3	-7.8	$ 1,225,125	4
Curt Casali	0.0	4.2	-4.2	$ 1,405,220	5
Francisco Cervelli	-0.3	0.2	-0.5	$ 7,300,000	5
Tom Murphy	-0.7	0.5	-1.2	$ 1,225,000	2
Cam Gallagher	-1.1	2.4	-3.5	$ 634,605	5
Kyle Higashioka	-1.1	3.3	-4.5	$ 702,733	3
Gary Sanchez	-1.4	-1.8	0.5	$ 3,699,683	6
Jonah Heim	-1.6	5.3	-6.9	$ 617,167	3
Geovany Soto	-2.3	-2.0	-0.3	$ 2,400,000	2
Carlos Ruiz	-2.4	-0.5	-1.9	$ 6,500,000	2
Max Stassi	-2.4	4.8	-7.2	$ 1,309,440	5
Ryan Jeffers	-2.5	4.0	-6.5	$ 644,275	2
Chad Wallach	-2.6	0.0	-2.6	$ 558,719	4

Travis d'Arnaud	-2.7	-1.2	-1.6	$ 4,988,767	6
Payton Henry	-2.8	-1.0	-1.8	$ 702,000	1
Francisco Mejia	-2.9	0.3	-3.2	$ 790,400	4
Raffy Lopez	-2.9	5.0	-7.9	$ 580,000	1
Austin Nola	-3.0	-3.0	0.0	$ 623,067	3
Jason Castro	-3.7	0.0	-3.7	$ 6,371,429	7
Carson Kelly	-3.8	0.0	-3.8	$ 1,541,675	4
Chris Herrmann	-3.9	-1.3	-2.6	$ 938,350	4
Bryan Holaday	-4.0	5.0	-9.0	$ 759,500	2
Rene Rivera	-4.0	0.0	-4.0	$ 2,275,000	2
Luke Maile	-4.2	2.3	-6.6	$ 679,133	3
Dustin Garneau	-4.7	0.0	-4.7	$ 593,500	2
Christian Vazquez	-4.8	5.3	-10.1	$ 3,257,000	7
Austin Wynns	-4.9	1.0	-5.9	$ 558,000	1
John Ryan Murphy	-5.1	2.3	-7.5	$ 662,233	3
Anthony Bemboom	-5.2	0.0	-5.2	$ 780,000	1
Steve Clevenger	-5.2	-2.0	-3.2	$ 516,500	1
Austin Hedges	-5.2	11.8	-17.0	$ 2,238,367	6
Reese McGuire	-5.4	1.5	-6.9	$ 645,900	2
Christian Bethancourt	-5.5	4.0	-9.5	$ 511,200	1
Bruce Maxwell	-5.6	0.0	-5.6	$ 550,000	1
Robinson Chirinos	-5.7	-3.7	-2.0	$ 3,025,000	6
Yadier Molina	-5.7	2.9	-8.5	$ 15,342,857	7
Eric Fryer	-5.8	0.0	-5.8	$ 625,000	2
Yan Gomes	-5.9	2.0	-7.9	$ 5,183,333	7
Eric Haase	-5.9	-9.0	3.1	$ 710,400	1
Omar Narvaez	-6.2	-5.5	-0.7	$ 1,984,367	6
Jeff Mathis	-6.2	7.8	-14.0	$ 2,350,000	5
Wilson Ramos	-6.2	-5.5	-0.7	$ 6,725,000	6
Sandy Leon	-6.7	8.7	-15.4	$ 1,750,000	3
Brian McCann	-7.2	-3.0	-4.2	$ 13,250,000	4
Tomas Nido	-7.2	6.0	-13.2	$ 674,673	3
Kevin Plawecki	-7.3	-1.7	-5.7	$ 1,159,389	6
Carlos Perez	-7.4	10.0	-17.4	$ 513,000	1
Jorge Alfaro	-7.9	-2.2	-5.7	$ 1,297,400	5
Victor Caratini	-8.0	-0.8	-7.3	$ 1,115,375	4
James McCann	-8.2	0.7	-8.9	$ 3,949,514	7
Chance Sisco	-8.2	-2.7	-5.6	$ 566,000	3
Devin Mesoraco	-8.3	-3.3	-4.9	$ 8,491,667	3
Austin Romine	-8.5	-0.8	-7.7	$ 1,651,833	6
Tony Wolters	-8.6	5.0	-13.6	$ 891,500	5
Caleb Joseph	-8.6	4.3	-12.9	$ 824,500	3
Jesus Sucre	-8.6	0.7	-9.3	$ 801,667	3
Chris Gimenez	-8.9	-3.5	-5.4	$ 962,500	2

Derek Norris	-8.9	7.5	-16.4	$ 2,062,500	2
Michael Perez	-9.0	1.0	-10.0	$ 571,600	3
Blake Swihart	-9.0	-1.0	-8.0	$ 563,500	1
Martin Maldonado	-9.0	5.3	-14.3	$ 3,035,714	7
Tucker Barnhart	-9.1	2.1	-11.2	$ 3,441,786	7
Brett Nicholas	-9.2	-4.0	-5.2	$ 537,000	1
Keibert Ruiz	-9.2	-4.0	-5.2	$ 701,300	1
Miguel Montero	-9.4	-1.5	-7.9	$ 14,000,000	2
Jose Lobaton	-9.4	-1.0	-8.4	$ 1,481,250	2
Grayson Greiner	-9.4	-2.3	-7.1	$ 571,167	3
Kevan Smith	-9.5	-6.5	-3.0	$ 566,250	2
Jett Bandy	-9.7	-3.5	-6.2	$ 545,250	2
Joey Bart	-9.8	-5.0	-4.8	$ 707,500	1
Aramis Garcia	-10.2	-1.5	-8.7	$ 637,750	2
Andrew Knapp	-10.3	-5.2	-5.1	$ 709,667	6
Kurt Suzuki	-10.3	-7.9	-2.4	$ 3,464,286	7
A.J. Ellis	-10.3	-5.0	-5.3	$ 2,750,000	3
Jose Herrera	-10.4	-1.0	-9.4	$ 700,000	1
Dom Nunez	-10.5	-4.0	-6.5	$ 638,750	2
Willians Astudillo	-10.5	-2.0	-8.5	$ 560,000	1
Erik Kratz	-10.6	0.5	-11.1	$ 1,175,000	2
Josh Phegley	-10.7	-3.8	-6.9	$ 759,375	4
Welington Castillo	-10.7	-5.8	-4.9	$ 6,050,000	4
Hank Conger	-11.1	-2.0	-9.1	$ 1,500,000	1
Stephen Vogt	-11.1	-3.6	-7.5	$ 2,068,500	5
Josh Thole	-11.2	1.0	-12.2	$ 800,000	1
Luis Torrens	-11.3	-5.5	-5.8	$ 721,125	4
Hector Sanchez	-11.5	-5.0	-6.5	$ 750,000	1
Alex Jackson	-11.8	0.0	-11.8	$ 580,500	1
Matt Wieters	-12.0	-2.8	-9.2	$ 8,060,000	5
Chris Iannetta	-12.4	-6.8	-5.6	$ 3,375,000	4
Pedro Severino	-12.4	-6.7	-5.7	$ 988,167	3
Chris Stewart	-12.4	-2.5	-9.9	$ 1,375,000	2
Trevor Brown	-12.5	-5.0	-7.5	$ 508,000	1
Ryan Hanigan	-12.9	-2.0	-10.9	$ 3,700,000	1
Drew Butera	-14.4	-9.0	-5.4	$ 1,654,167	3
Riley Adams	-15.1	-6.0	-9.1	$ 706,700	1
Andrew Knizner	-16.3	-5.0	-11.3	$ 648,750	2
Cameron Rupp	-16.5	-10.0	-6.5	$ 541,500	2
John Hicks	-17.0	-3.5	-13.5	$ 559,200	2
Elias Diaz	-17.1	-7.0	-10.1	$ 1,334,375	4
Nick Hundley	-17.4	-10.0	-7.4	$ 2,225,000	4
Stuart Turner	-17.4	-5.0	-12.4	$ 535,000	1
Jarrod Saltalamacchia	-17.8	-8.0	-9.8	$ 507,500	1

Jonathan Lucroy	-18.5	-9.5	-9.0	$ 4,862,500	4
Francisco Pena	-18.6	-6.0	-12.6	$ 650,000	1
Zack Collins	-22.3	-18.0	-4.3	$ 575,000	1
A.J. Pierzynski	-28.5	-7.0	-21.5	$ 3,000,000	1
Dioner Navarro	-33.4	-15.0	-18.4	$ 4,000,000	1

E. Team Catcher Rankings

App. Table 10
Catcher Rankings by Team

This table lists the all the teams in the sample that includes team mandated players. Avg. total performance is the average of the sum of all the catchers' average DRS and average OFF statistics who played for that team between 2016-2022. Average salary is the average salary of all the catcher who played for that team over the period of analysis (2016-2022). Likewise, # of catchers is how many catchers played for that team between 2016-2022.

Team	Avg. Tot. Perf.	Avg. DRS	Avg. OFF	Avg. Salary	# Of Catchers
LAD	9.7	6.9	2.8	$ 3,474,154	13
CHC	3.2	2.3	0.8	$ 3,981,625	12
SFG	1.3	1.5	-0.2	$ 11,670,439	10
ATL	1.3	1.2	0.1	$ 3,977,692	13
TOR	0.3	3.5	-3.3	$ 4,806,315	13
MIL	-0.5	3.5	-4.0	$ 2,818,369	13
PHI	-0.6	-2.1	1.5	$ 5,043,654	13
SEA	-1.5	-1.2	-0.3	$ 1,635,525	12
KCR	-1.9	-0.7	-1.3	$ 5,164,271	13
NYY	-2.0	0.9	-3.0	$ 2,663,379	14
ARI	-2.9	3.3	-6.2	$ 1,921,873	15
MIN	-3.4	-1.1	-2.2	$ 3,379,650	15
MIA	-3.9	-1.3	-2.6	$ 1,224,805	16
PIT	-3.9	1.1	-5.0	$ 2,910,792	12

TBR	-4.5	1.9	-6.4	$ 2,131,421	14
CHW	-4.7	-4.4	-0.3	$ 6,452,692	13
CIN	-4.9	1.3	-6.2	$ 2,420,036	14
HOU	-5.4	0.2	-5.6	$ 6,018,827	11
CLE	-5.5	5.5	-11.0	$ 2,687,220	15
TEX	-5.9	0.8	-6.7	$ 1,664,500	11
BOS	-6.2	3.8	-9.9	$ 2,151,042	12
LAA	-6.4	1.8	-8.2	$ 1,584,500	9
OAK	-6.6	-2.1	-4.6	$ 1,114,533	15
SDP	-7.4	2.1	-9.6	$ 1,076,727	15
STL	-8.2	0.2	-8.5	$ 8,732,500	13
- - -	-8.3	-1.9	-6.4	$ 3,726,228	40
WSN	-8.7	-3.1	-5.6	$ 4,441,986	12
NYM	-8.9	0.1	-8.9	$ 3,664,587	11
BAL	-9.4	-2.6	-6.8	$ 2,245,964	14
COL	-13.2	-2.0	-11.2	$ 1,736,250	12
DET	-14.3	-3.5	-10.8	$ 1,595,533	12

F. Baseball Statistics Definitions

App. Table 11
Baseball Statistics Definitions

This table provides a brief definition of statistics mentioned in the paper. All definitions come directly from the official MLB Glossary (“Glossary” 2022).

Variable Name	Definition
Batting average (BA)	A player's hits divided by his total at-bats. Reported as a number between zero (.000) and one (1.000).
Runs batted in (RBI)	A batter is credited with an RBI in most cases where the result of his plate appearance is a run being scored.
Earned runs average (ERA)	The number of earned runs a pitcher allows per nine innings. Earned runs are any runs that scored without the aid of an error or a passed ball.
On base percentage (OBP)	How frequently a batter reaches base per plate appearance. Times on base include hits, walks and hit-by-pitches, but do not include errors, times reached on a fielder's choice or a dropped third strike.
Weighted on base average (wOBA)	A version of on-base percentage that accounts for how a player reached base -- instead of simply considering whether a player reached base.
Slugging percentage (SLG)	The total number of bases a player records per at-bat. Unlike on-base percentage, slugging percentage deals only with hits and does not include walks and hit-by-pitches in its equation.

9. Bibliography

"2017-2021 Collective Bargaining Agreement." *Mlbplayers.com*, MLBPA, 2016, https://www.mlbplayers.com/_files/ugd/b0a4c2_95883690627349e0a5203f61b93715b5.pdf.

Adams, Steve. "Details on Bonus Pool for Pre-Arbitration Players." *MLB Trade Rumors*, 9 Sept. 2022, https://www.mlbtraderumors.com/2022/09/mlb-pre-arbitration-bonus-pool-payouts.html.

Andrecheck, Sky. "WAR, Salary, and Service: Estimating Dollars per Win." *The Baseball Analysts*, 2009, http://baseballanalysts.com/archives/2009/11/war_salary_and.php.

Avidon, Eric. "Baseball's Rangers Seek Analytics Edge with Tableau." *Techtarget.com*, 2022, https://www.techtarget.com/searchbusinessanalytics/feature/Baseballs-Rangers-seek-analytics-edge-with-Tableau.

Beneventano, Philip, et al. "Predicting Run Production and Run Prevention in Baseball: The Impact of Sabermetrics." *Researchgate.net*, 2012, https://www.researchgate.net/profile/Bruce-Weinberg/publication/266344641_Predicting_Run_Production_and_Run_Prevention_in_Baseball_The_Impact_of_Sabermetrics/links/54bab6740cf253b50e2d0563/Predicting-Run-Production-and-Run-Prevention-in-Baseball-The-Impact-of-Sabermetrics.pdf?origin=publication_detail.

Blass, Asher A. "Does the Baseball Labor Market Contradict the Human Capital Model of Investment." *Jstor.org*, 1992, https://www.jstor.org/stable/2109657.

Brown, Daniel T., et al. "Moneyball After 10 Years: How Have Major League Baseball Salaries Adjusted." *Sage Journals*, Oct. 2015, https://journals.sagepub.com/doi/abs/10.1177/1527002515609665?casa_token=FK5kc0rK6KIAAAAA%3AQN4yJu2k9Cm-hdwqh89nyZ2spjBvsUcUubafWVSUIjA_udjH91a0-NLIytWiks4Zlyzy9OtE8krzNw&.

Clayton, Matthew, and David Yermack. "Major League Baseball Player Contracts: An Investigation of the Empirical Properties of Real Options." *SSRN*, 3 Nov. 2001, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1294593.

Clemens, Ben. "The Count Is King (Even after Accounting for Batter Skill)." *FanGraphs Baseball*, 8 Jan. 2020, https://blogs.fangraphs.com/the-count-is-king-even-after-accounting-for-batter-skill/.

DaSilva, Alex. “Catcher Framing Seasonal Summary.” *RPubs*, 2021, https://rpubs.com/alexdasilva/CatcherFramingSeasonal.

Dewan, John, and Ben Jedlovec. “The Book--Playing the Percentages in Baseball.” *THE BOOK*, 2009, http://www.insidethebook.com/ee/index.php/site/comments/john_dewan_and_research_assistant_speak/#3.

Dewan, John, et al. “Catcher Stolen Base Runs Saved.” *Fielding Bible*, 2020, https://fieldingbible.com/methodology-catchersb.asp.

Dewan, John, et al. “Frequently Asked Questions about Plus/Minus and Runs Saved.” *Fielding Bible*, 2020, https://fieldingbible.com/Fielding-Bible-FAQ.asp.

Dewan, John, et al. “Methodology: Adjusted Earned Runs Saved.” *Fielding Bible*, 2020, https://fieldingbible.com/methodology-adjer.asp.

Dewan, John, et al. “Methodology: Good Play/Misplay Runs Saved.” *Fielding Bible*, 2020, https://fieldingbible.com/methodology-gfpdme.asp.

Dewan, John, et al. “Methodology: Plus/Minus.” *Fielding Bible*, 2020, https://fieldingbible.com/methodology-plusminus.asp.

Domazlicky, Bruce R, and Peter M Kerr. “Baseball Attendance and the Designated Hitter.” *Sage Journals*, Mar. 1990, https://journals.sagepub.com/doi/10.1177/056943459003400107.

Ehrlich, Justin Andrew, and Joel M. Potter. “Is Offense Worth More than Defense and Pitching? Marginal Revenue Product and Revenue Sharing in Major League Baseball.” *Managerial Finance*, Emerald Publishing Limited, 4 July 2020, https://www.emerald.com/insight/content/doi/10.1108/MF-04-2020-0210/full/html.

Estes, Brent C. “Predicting Productivity in a Complex Labor Market: A Sabermetric Assessment of Free Agency on Major League Baseball Player Performance.” *The Case Centre*, 2006, https://www.thecasecentre.org/products/view?id=126510.

Fast, Mike. “Spinning Yarn: Removing the Mask Encore Presentation.” *Baseball Prospectus*, 24 Sept. 2011, https://www.baseballprospectus.com/news/article/15093/spinning-yarn-removing-the-mask-encore-presentation/.

Feinsand, Mark. “MLB, MLBPA Agree to New CBA; Season to Start April 7.” *MLB.com*, MLB, 11 Mar. 2022, https://www.mlb.com/news/mlb-mlbpa-agree-to-cba.

Feinsand, Mark. "Play Ball: MLB Announces 2020 Regular Season." *MLB.com*, MLB, 6 July 2020, https://www.mlb.com/news/mlb-announces-2020-regular-season.

Garro, Adrian, and Eric Chesterson. "Which Is the More Demanding Job: Pitcher or Catcher?" *MLB.com*, 11 Feb. 2019, https://www.mlb.com/cut4/which-is-harder-being-a-catcher-or-a-pitcher-c303719102#:~:text=But%20the%20catcher%20has%20the,the%20hardest%20job%20in%20baseball.

"Glossary." *MLB.com*, Major League Baseball, 2022, https://www.mlb.com/glossary.

Hakes, Jahn K., and Raymond D. Sauer. "An Economic Evaluation of the Moneyball Hypothesis." *Jstor.org*, American Economic Association, 2006, https://www.jstor.org/stable/pdf/30033672.pdf.

Hakes, Jahn K., and Raymond D. Sauer. "The Moneyball Anomaly and Payroll Efficiency: A Further Investigation." *Researchgate.net*, International Journal of Sport Finance, Nov. 2007, https://www.researchgate.net/publication/5142757_The_Moneyball_Anomaly_and_Payroll_Efficiency_A_Further_Investigation.

Healey, Glenn. "The New Moneyball: How Ballpark Sensors Are Changing Baseball." *IEEE Xplore*, 2017, https://ieeexplore.ieee.org/abstract/document/8074554.

Hook, Christian. "Credited Framing Runs: Which Catchers Saved the Most Runs with Their Receiving?" *Medium*, Medium, 2 July 2020, https://medium.com/@ChristianHook/credited-framing-runs-which-catchers-saved-the-most-runs-with-their-receiving-e82ee8e9c546.

Humphreys, Michael. "Wizardry: Baseball's All-Time Greatest Fielders Revealed." *ProQuest Ebook Central*, 2011, https://ebookcentral.proquest.com/lib/claremont/reader.action?docID=679339.

Kelly, Matt. "What Is Sabermetrics?" *MLB.com*, MLB, 27 May 2019, https://www.mlb.com/news/sabermetrics-in-baseball-a-casual-fans-guide.

Koprivica, Bojan. "Another One Bites the Dust." *The Hardball Times*, 2011, https://tht.fangraphs.com/another-one-bites-the-dust/.

Kuty, Brendan. "The Special Insight the Yankees Had into New Catching Whiz." *Silive*, 7 Apr. 2022, https://www.silive.com/yankees/2022/04/the-special-insight-the-yankees-had-into-new-catching-whiz.html?outputType=amp.

Lindbergh, Ben. "Yadi's Secret Sauce." *" Yadi's Secret Sauce*, 2014, http://grantland.com/features/st-louis-cardinals-yadier-molina-catcher-value-game-calling/.

Loughin, Thomas M, and Jason L Bargen. "Assessing Pitcher and Catcher Influences on Base Stealing in Major League Baseball." *Taylor & Francis*, 2008, https://www.tandfonline.com/doi/abs/10.1080/02640410701287255.

Magel, Rhonda, and Michael Hoffman. "Predicting Salaries of Major League Baseball Players." *International Journal of Sports Science*, Scientific & Academic Publishing, 2015, http://article.sapub.org/10.5923.j.sports.20150502.02.html.

"Major League Baseball Signing Bonus Research." *Baseball Almanac*, 2022, https://www.baseball-almanac.com/players/baseball_signing_bonus.shtml#:~:text=A%20baseball%20signing%20bonus%20or,and%20begin%20their%20professional%20careers.

Marchi, Max. "Evaluating Catchers: Quantifying the Framing Pitches Skill." *The Hardball Times*, 10 June 2011, https://tht.fangraphs.com/evaluating-catchers-quantifying-the-framing-pitches-skill/.

Maurice, Chris. "The Value of Major League Baseball Players." *Brynmar.edu*, 2010, https://scholarship.tricolib.brynmawr.edu/bitstream/handle/10066/5985/2010MauriceC_Thesis.pdf?sequence=1&isAllowed=y.

Mizels, Joshua, et al. "Current State of Data and Analytics Research in Baseball." *SpringerLink*, Springer US, 29 Apr. 2022, https://link.springer.com/article/10.1007/s12178-022-09763-6.

"MLB Positional Payrolls." *Spotrac.com*, 2022, https://www.spotrac.com/mlb/positional/.

Pemstein, Jonah. "Introducing the Batter-Specific Run-Expectancy Tool." *FanGraphs Baseball*, 2016, https://blogs.fangraphs.com/introducing-the-batter-specific-run-expectancy-tool/.

"Plate Appearance (PA): Glossary." *MLB.com*, 2022, https://www.mlb.com/glossary/standard-stats/plate-appearance.

Pollack, Brian. "What Gets Paid? Analyzing the Major League Baseball Contract Market." *Dukespace*, 2017, https://dukespace.lib.duke.edu/dspace/bitstream/handle/10161/14325/Pollack2017.pdf?sequence=.

Rosales, Joe, and Scott Spratt. "Who Is Responsible for a Called Strike?" *MIT Sloan Sports Analytics Conference*, 2015, https://www.sloansportsconference.com/research-papers/who-is-responsible-for-a-called-strike.

"Salary Arbitration: Glossary." *MLB.com*, 2022, https://www.mlb.com/glossary/transactions/salary-arbitration.

"Service Time: Glossary." *MLB.com*, 2022, https://www.mlb.com/glossary/transactions/service-time.

Slowinski, Piper. "DRS." *Sabermetrics Library*, 2010, https://library.fangraphs.com/defense/drs/.

Slowinski, Piper. "WOBA." *Sabermetrics Library*, 2010, https://library.fangraphs.com/offense/woba/.

Slowinski, Piper. "WRAA." *Sabermetrics Library*, 2010, https://library.fangraphs.com/offense/wraa/.

Smith, Erin E, and Jon D Groetzinger. "Do Fans Matter? the Effect of Attendance on the Outcomes of Major League Baseball Games." *De Gruyter*, De Gruyter, 28 Jan. 2010, https://www.degruyter.com/document/doi/10.2202/1559-0410.1192/html.

Swartz, Matt. "Ahead in the Count: Methodology of the New Morp." *Baseball Prospectus*, 23 Apr. 2010, https://www.baseballprospectus.com/news/article/10642/ahead-in-the-count-methodology-of-the-new-morp/.

Tainsky, Scott, and Jason A. Winfree. "Short-Run Demand and Uncertainty of Outcome in Major League Baseball." *SpringerLink*, Springer US, 23 Oct. 2010, https://link.springer.com/article/10.1007/s11151-010-9267-3.

Wasserman, Tyler. "Determinants of Major League Baseball Player Salaries." *SURFACE at Syracuse University*, May 2013, https://surface.syr.edu/honors_capstone/99/.

Weinberg, Neil. "Getting Started." *Sabermetrics Library*, 2015, https://library.fangraphs.com/getting-started/.

Weinberg, Neil. "Off." *Sabermetrics Library*, 2014, https://library.fangraphs.com/offense/off/.

Weinberg, Neil. "RE24." *Sabermetrics Library*, 2014, https://library.fangraphs.com/misc/re24/.

Weinberg, Neil. "What Do We Know about Catcher Defense?" *Sabermetrics Library*, 2014, https://library.fangraphs.com/what-do-we-know-about-catcher-defense/.

Weinstein, Max. "How Do We Assign Credit for Catching Base-Stealers?" *The Hardball Times*, 18 July 2014, https://tht.fangraphs.com/how-do-we-assign-credit-for-catching-base-stealers/.

Williams, Jack F. "The Coming Revenue Revolution in Sports." *Willamette.edu*, 2006, https://willamette.edu/law/pdf/review/42-4/williams.pdf.

"WOBA & FIP Constants." *Guts! | FanGraphs Baseball*, 2022, https://www.fangraphs.com/guts.aspx?type=cn.