Labor Market Efficiency and Team Transactions in Major League Baseball

by

Ryan Xavier Pinheiro

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Sport Management)
in The University of Michigan
2022

Doctoral Committee:

Professor Stefan Szymanski, Chair
Professor Charlie Brown
Professor Rod Fort
Professor Brad Humphreys, West Virginia University

Ryan Xavier Pinheiro

rxp@umich.edu

ORCID iD: 0000-0002-3614-7392

ACKNOWLEDGEMENTS

Thanks to my advisor and my other dissertation committee members for their help, support, expertise, and guidance over the past few years. Also thanks to my parents, sister, extended family, and friends for their love and support throughout my time at Michigan. Finally, I would like to thank God for an amazing experience at the University of Michigan and for all the people that have helped me along the way.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

This dissertation is comprised of three separate papers related to the player labor market and team decision making process in Major League Baseball (MLB).

The first paper focuses on the MLB labor market after the publication of the book *Moneyball* (Lewis, 2003). The book claimed that data analytics enabled savvy operators to exploit inefficiencies in the market for baseball players. The economic analysis of Hakes and Sauer (2006) appeared to show that the publication of *Moneyball* represented a watershed, after which inefficiencies had been competed away. In both cases analysis focused on composite statistics such as on-base percentage (OBP) and slugging percentage (SLG). This paper relies on a more structural approach, associated with the statistical analysis of Lindsey (1963) which identifies the run value of each individual event in a game. Using a dataset of every event in every game from 1996 to 2015, we show that the run value of each event can be accurately calculated, as can the run value contribution of each player. We show that the compensation of free agents reliably reflects the run value contribution of each player, regardless of the source of those contributions (walks, singles, and home runs). We find this was true both before and after the publication of *Moneyball*, suggesting that the labor market for batters in Major League Baseball operated efficiently across our entire sample period.

The second paper proposes novel approaches to measuring team productivity and evaluating trading efficiency in MLB from 1994-2016 through an application of portfolio theory.

The performance of individual players is again measured using a structural approach relating player outcomes to team runs as developed by Lindsey (1963). Using a portfolio theory framework, we treat MLB teams as a portfolio of players (assets), each of which can be defined by an expected contribution of runs per game and the variance of this measure. It is found that both the expected value and variance have a positive impact on team runs scored. Given our definition of teams characterized by their expected values and variances, we evaluate trading efficiency between teams given their pre-trade expected values and variances and the acquired player's pre-trade expected value and variance. We find that trade efficiency has improved in recent years, consistent with the growth in data-driven decision making used in MLB front offices.

The third and final paper investigates the impact of employer role assignments for entry-level positions on the career paths of their entry-level employees. Specifically, we focus on the impact of a pitcher's entry-level role on the pitcher's survivability in MLB. In professional baseball, pitchers are generally categorized into one of two pitcher types: starting pitcher or relief pitcher. These pitcher types have distinct roles and different levels of labor supply associated with them. These roles are also generally assigned to the pitcher by his team. Using Kaplan-Meier estimators and a Cox-Proportional Hazard model, we find that players assigned entry-level reliever roles face a significant survival disadvantage compared to entry-level starting pitchers and that differences in quality between pitcher groups do not explain this disadvantage. These results illustrate how the early-career decisions made by employers on behalf of their employees can ultimately have career altering implications for workers.

Chapter I

Introduction

This dissertation is comprised of three separate papers related to the player labor market and team decision making process in Major League Baseball (MLB). In recent years, there has been a substantial growth of academic research in the areas of sport economics and sport finance. One key area of research within these disciplines has been on the labor market of professional athletes. As alluded to in Kahn (2000), the professional sports industry provides researchers a unique opportunity to test theory from labor economics which may be otherwise difficult when conducting such research in other industries. Desirable characteristics of using professional sports data for labor economics research include the observability of employee performance, the ability to objectively measure employee productivity, and the availability of employee wage data.

Baseball has been one of the more popular sports for research in labor economics due to the relative ease in reliably measuring a player's contributions to his team. A baseball game is comprised of a series of largely discrete events for which production value can be largely attributed to two primary players: the pitcher and the hitter. This makes performance measurement in baseball easier compared to other sports that are more continuous and team-dependent in nature (i.e. soccer, basketball, etc.).

This dissertation proceeds as follows. Chapter 2 focuses on the MLB labor market after the publication of the book *Moneyball* (Lewis, 2003). This chapter seeks to determine whether significant changes to the salary determination process exist for free agents after the publication of *Moneyball* using a more structural approach to productivity measurement than the previous research on the topic. Chapter 3 develops an approach utilizing the mean-variance framework from financial portfolio theory to measure trade efficiency between teams in MLB. Chapter 4 investigates the impact of entry-level roles on career paths in MLB. This chapter specifically focuses on the entry-level roles of pitchers and how these entry-level role assignments greatly impact the ability of pitchers to survive in the league long-term. Chapter 5 offers a brief conclusion of this dissertation with a high-level summary of results from each chapter.

All Runs Are Created Equal: Labor Market Efficiency in Major League Baseball

2.1   Introduction

"Every event on a baseball field Paul understood as having an "expected run value." You don't need to be able to calculate expected run values to understand them. Everything that happens on a baseball field alters, often very subtly, a team's chances of scoring runs. Every event on a baseball field changes, often imperceptibly, the state of the game. For example, the value of having no runners on base with nobody on base and no count on the batter is roughly .55 runs, because that is what a baseball team, on average, will score in that situation. If the batter smacks a double, he changes the "state" of the game: it's now nobody out with a runner on second base. The expected run value of that new "state" is 1.1 runs. It follows that the contribution of a leadoff double to a team's expected runs is .55 runs (1.1 minus .55). If the batter, instead of hitting a double, strikes out, he lowers the team's expected run value to roughly .30 runs. The cost of making that out was therefore .25 runs — the difference between the value of the original state of the game and the state the batter left it in." *Moneyball* (Lewis, 2003, p. 134)

"This is a very simple game. You throw the ball, you catch the ball, you hit the ball." Bull Durham (1988)

Michael Lewis's book *Moneyball: the Art of Winning an Unfair Game* is widely credited with transforming the analysis of baseball. The central proposition of the book is that

general managers and coaches used to neglect statistics and rely on gut-feel or the opinions of scouts who paid little attention to performance data. According to the book, Billy Beane, general manager of the Oakland Athletics and Paul DePodesta (mentioned in the quote above) generated above average team performance on a restricted budget by hiring players with undervalued statistics, such as the capacity to draw a walk.

Hakes and Sauer (2006, 2007) appeared to confirm this finding. Their research focused on two statistics – slugging percentage (SLG) and on-base percentage (OBP). The former, a measure of batting power, was popular as an index of player ability but gave no weight to the capacity to draw a walk. By contrast, on-base percentage was less popular as a measure of ability but crucially did include the capacity to draw a walk. The authors compared the significance of the two statistics in (a) determining team wins and (b) determining player salaries. Efficiency, they argued, required that the relative weights of each statistic in a wins regression should be similar to the corresponding weights in a salary regression. In their regression analysis they found the weights were relatively stable, year by year, in the wins regression. In the salary regression, however, they found that prior to publication of *Moneyball*, SLG contributed significantly, but that OBP was statistically insignificant. After the publication of *Moneyball*, however, OBP did become significant in the salary regression and the relative size of the OBP and SLG coefficients were similar to what was observed in the wins regression. This, they argued, provided evidence that the publication of *Moneyball* really did change the relative valuation of skills and made the player market more efficient.

The book frequently mentioned the undervaluation of OBP, and the striking confirmation in the work of Hakes and Sauer has emphasized this point further. After the publication of *Moneyball*, both OBP and SLG became frequently referenced performance indicators in Major League Baseball as both metrics improved upon traditional measures (such as batting average) by aggregating and distinguishing between more types of player performance outcomes than previously used metrics. But as the quotation at the start of this paper

4

suggests, DePodesta and Beane were ultimately interested in player performance outcomes at a more granular level than OBP and SLG. In this paper we develop an approach based around the concept of run expectancy, which allows us to structurally determine the value of specific events in baseball, based on each event's contribution to runs. While this approach is popular in the sabermetrics community, it has yet to be utilized extensively in the academic research on labor market efficiency in Major League Baseball. Given the interest in more granular levels of player performance as evidenced by the quote above, we feel that a logical next step is to fill this gap in the literature. We follow the approach of Hakes and Sauer (2006), measuring the determinants of win percentage at the level of the team and then testing whether the contributions of players are rewarded in proportion to their contributions to win percentage. To do so, however, we adopt the concept of run expectancy and the run value produced by players.

If the player market works efficiently, then the batting contributions of players should be valued purely on the basis of their contributions to generating runs, and not the way in which those runs are generated. Thus walks, singles, and home runs are all batting events that contribute to run scoring. They each generate different amounts of expected runs, which we can calculate in any given season by the frequency of run changes associated with each event type, based on all events in a season (of which there are around 200,000). Thus for example, if the expected run value of a home run was 1.5 runs, and the expected run value of a single was 0.5 runs, then the player who hits 12 singles contributes exactly as much run value as a player who contributes 4 home runs (all else equal).

In this paper we derive the run value contributions of Major League Baseball players and test the proposition that batting event types are valued proportionately to their run values, using player performance data covering the period 1996-2015 and salaries covering the period 1997-2016. When considering free agents we find that we cannot reject the hypothesis of equal returns across the entire period, and find only small differences when comparing the

periods before and after the publication of *Moneyball*. We do observe an increase in the return to walks compared to singles and home runs, but the increase does not lead to significantly different returns for these outcomes.

In the next section we review the literature, and the following section we describe our model and estimation. The final section concludes.

2.2   Literature Review

Economists have long had an interest in productivity and salary determination in baseball. Scully (1974) identified the relationship between (a) batting performance and wins (marginal product) and (b) wins and team revenue (marginal revenue) in order to identify marginal revenue product and measure the extent of monopsonistic exploitation under the reserve clause. The introduction of limited free agency after 1975 led to follow up studies by Sommers and Quinton (1982), Raimondo (1983), Fort (1992), Zimbalist (1992), Kahn (1993), MacDonald and Reynolds (1994), Marburger (1994), Krautmann (1999), Hakes and Turner (2011), Bradbury (2013) and Humphreys and Pyun (2017). A parallel development during this period was the development of "sabermetrics" – the analysis of baseball statistics. The most notable figure in this movement has been Bill James, who developed a number of statistics for measuring player performance, popularized the use of statistical methods and published numerous books on baseball statistics (e.g. James (2003)). He, in turn, was a significant influence on Billy Beane, who became the general manager of the Oakland A's in the 1990s and whose baseball strategy was the subject of *Moneyball*.

The papers of Hakes and Sauer (2006, 2007) both formalized some propositions contained in *Moneyball* and provided a statistical test. Thus, the book refers to DePodesta's views on "the overwhelming importance of on-base percentage" (Lewis, 2003, p34) and states that "not long after he arrived in Oakland, Paul asked himself a question: what was the

relative importance of on-base percentage and slugging percentage" (Lewis, 2003, p127). That was exactly the focus of Hakes and Sauer (2006). They showed that OBP has been consistently important in determining team win percentage, at least as important as SLG, but that only after the publication of *Moneyball* did it become significant in determining player salaries (whereas prior to Moneyball only SLG had been significant). Hakes and Sauer (2007) extended the analysis one year further forward and fifteen years further back, finding that, while a trend increase in the valuation of OBP pre-dated *Moneyball*, there was still a significant jump in the three years after publication. They also develop alternative measures, intended to capture the notions of hitting ability, control (which includes the capacity to draw a walk), and power to better isolate the skills underlying OBP and SLG.

SLG is a weighted measure of batting performance, equal to (singles + 2 x doubles + 3 x triples + 4 x home runs)/at bats (despite the name, it is not a percentage). OBP is an unweighted measure of the capacity to get on base, either by hitting the ball or drawing a walk. Thus OBP = (singles + doubles + triples + home runs + walks + hit by pitch)/(at bats + walks + hit by pitch + sacrifice flys), and hence, is a percentage.[1] As these definitions might suggest, there is in fact a strong correlation between OBP and SLG, a point which is made by Hakes and Sauer (2007), Baumer and Zimbalist (2014), and Holmes et al (2018). The fact that the collinearity here is partly structural, i.e. caused by designing statistics that include the same measures, suggests that we can address the problem by relying on more disaggregated statistics that are not as closely correlated. This approach is adopted by Duquette et al (2019), who use elements of team batting performance (singles, doubles, triples, home runs, walks), base running (stolen bases and caught stealing), pitching performance (strikeouts) and fielding (errors) to explain both win percentage and team

---

[1]For those not familiar with baseball, a batter that is struck by a ball thrown by the pitcher gets to walk to first base, while a sacrifice fly is a ball into the air leading to the player being caught out, but allows a player already on base to score a run (reaching home plate). When a batter goes in to bat this is called a plate appearance, but an at bat refers only to those plate appearances which do not result in a walk, hit by pitch or sacrifice fly.

payrolls. As we discuss below, we can in fact be much more precise in relating batting performance (inputs) to runs produced (output).

Baumer and Zimbalist (2014) conduct a follow up study to Hakes and Sauer using data for the seasons 1985-2012. In their salary regressions, they use the alternative metrics pertaining to hitting ability, control, and power developed in Hakes and Sauer (2007) in place of OBP and SLG. They find that the market inefficiencies detected by Hakes and Sauer had been corrected by 2006.

Brown et al (2017) extended the Hakes and Sauer analysis another five years further forward, up until 2011, eight years after the publication of *Moneyball*. They adopted the same methodology as Hakes and Sauer (2006), focusing on the impact of OBP and SLG. When calculating salary regressions they estimated separate regressions for the three contract states of players: reserve clause (one to two years service time in the majors), arbitration eligible (3-6 years) and free agents (more than six years).[2] They conclude that the step jump identified by Hakes and Sauer has in fact persisted, although they only find an effect for free agents.

Holmes et al (2018) use data for the period 1997-2012. They test a model similar to Hakes and Sauer, but find that OBP increased on salaries only in the year 2004, and not thereafter. However, they are also skeptical about this result given the correlation between OBP and SLG. Therefore, they also use the alternative measures of hitting ability, control and power developed in Hakes and Sauer (2007). However, they find little evidence with these measures of any appreciable increase in the salary returns to control and thus conclude that there is little evidence that anything changed in the post-*Moneyball* era.

Duquette et al (2019) use team payroll data from 1988 to 2017 to conduct their pre- and post-*Moneyball* analysis. Unlike the other studies, they take the more disaggregated

---

[2]Hakes and Sauer included all players in the same equation but identified contract status with dummy variables.

approach, and regress payroll on team level aggregates for walks, singles, doubles, triples, home runs, stolen bases, strikeouts and errors. They find that the coefficient for walks is significant before the publication of *Moneyball* and that changes to the coefficient after *Moneyball's* publication are insignificant. Unlike all of the other studies, Duquette et al (2019) do not use individual player salary data, thus combining players of all contractual types, ages, experience and so on. Since team payrolls in baseball are closely correlated with win percentage (see. e.g. Hall et al (2002)), these regressions may just be capturing the relationship between performance statistics and winning, rather than the valuation of the individual skills of each player.

The links between the batting statistics and outcomes described thus far are ad-hoc. The batting statistics are plausible contributors to team success, but the causal link from the statistics to runs and therefore wins is not explicit. A structural approach to the evaluation of batting performance has been developed in recent years, building on the pioneering work of Lindsey (1963). Lindsey formalized the notion of run expectancy. Expected runs are defined by the frequency of runs scored from a given state in a half inning to the end of the half inning.[3] There are 24 possible states in a half inning depending on the number of outs (0, 1 or 2) and the runners on base (no runners on base, a runner on first and no one on second and third, a runner on first and second and no one on third, etc). In the game, the states change as a result of events - most events are the result of an at-bat, e.g. a single, a home run, a walk or an out. Each event alters the expected number of runs that will be scored from the current state to end of the half inning. Lindsey calculated frequencies by hand based on 27,027 events over two seasons.

Thorn et al (1985) extended Lindsey's work and used computers to make these calculations on the basis of all events in a single season, so as to generate a run expectancy matrix

---

[3]In baseball, the inning refers to the at bats of both teams, and so the performance of just one team in an inning is referred to as a half inning. This is unlike cricket, where the inning refers typically to just a single team.

for a given year. Once the run expectancy matrix is derived, it is possible to place a run value on a specific event equal to the runs scored during the event plus the run expectancy at the end of the event minus the run expectancy at the beginning of the event. This method can then be used to value the run contributions of individual players. This approach has been extensively discussed in the work of Albert and Bennett (2001) and Tango et al (2006). Marchi and Albert (ch.5, 2013) provide a useful guide on how to generate these statistics, while Baumer et al (2015) show how to derive the related measure of Wins Above Replacement (WAR) using open-source data. Despite this previous research by the sabermetrics community, the run value approach to valuing player productivity has yet to be utilized extensively in the sport economic literature. While some economists have discussed the usefulness of run expectancy and related concepts (see e.g. Bradbury [2010], Baumer and Zimbalist [2014]), these tools have been relatively neglected in the field up until now. However, using the run value approach to measure player productivity allows for a more explicit determination of a player's contribution to team output than the measures of productivity included in past research. As a result, the run value approach should allow us to conduct a reliable test of market efficiency for the more granular outcomes underlying OBP and SLG. Thus, an application of the run value approach in a test for market efficiency seems like a natural next step in extending the previous work completed on player pay and productivity in baseball. In the following sections, we use these methods to test the efficiency of the labor market and to test for *Moneyball* effects in Major League Baseball.

2.3   Model and Estimation

Our approach essentially mimics that of Hakes and Sauer (2006). First, we estimate the determinants of win production using team level statistics, and then we estimate the return to players on their contributions to those same statistics. We test the hypothesis that the

returns to different types of event production are equalized, so that player event contributions are valued strictly in proportion to their contribution to winning. In our paper, we base the estimates on statistics that should translate directly into runs, which we take to be the key measure of batting performance in baseball.

2.3.1 Team Statistics

From the point of view of any team i we can define the outcome in any season t as a function of runs scored and runs against:

$$(2.1) \qquad\qquad Wpc_{it} = f(R_{it}, RA_{it})$$

where Wpc is win percentage, R is team runs scored and RA is team runs against.

We can then define runs as a function of batting events, which advance the players around the bases:

$$(2.2) \qquad\qquad R_{it} = g(BB_{it}, X1B_{it}, X2B_{it}, X3B_{it}, HR_{it})$$

where BB refers to walks, $X1B$ to singles, $X2B$ to doubles, $X3B$ to triples and HR to home runs. Likewise runs against can be defined as:

$$(2.3) \qquad\qquad RA_{it} = h(BBA_{it}, X1BA_{it}, X2BA_{it}, X3BA_{it}, HRA_{it})$$

where each total is analogous to (2.2), only referring to the scores against the team in question.

In Table 2.1 we report ordinary least squares estimates for the linear versions of equations

(2.1)-(2.3) using data aggregated at the team and season level for the seasons 1996-2015.[4] We obtain this team-level data from the Lahman Database.[5] These estimates are the analogue of calculating the relationship between Wpc and OBP and SLG in Hakes and Sauer (2006, Table 1). Team and season level fixed effects are included in these models as well.

Table 2.1: Team Performance Estimates

Dependent variable	Wpc	R	RA	Wpc	Wpc
R	0.000593***				
	(0.0000172)				
RA	−0.000623***				
	(0.0000155)				
BB		0.347***		0.000221***	0.000264***
		(0.0194)		(0.0000262)	(0.0000182)
X1B		0.498***		0.000330***	0.000330***
		(0.0196)		(0.0000263)	(0.0000193)
X2B		0.730***		0.000337***	0.000399***
		(0.0461)		(0.0000622)	(0.0000420)
X3B		1.402***		0.000419*	0.000690***
		(0.137)		(0.000185)	(0.000133)
HR		1.398***		0.000953***	0.000952***
		(0.0405)		(0.0000555)	(0.0000467)
BBA			0.378***	−0.000307***	−0.000264***
			(0.0215)	(0.0000268)	(0.0000182)
X1BA			0.534***	−0.000330***	−0.000330***
			(0.0230)	(0.0000285)	(0.0000193)
X2BA			0.739***	−0.000418***	−0.000399***
			(0.0455)	(0.0000564)	(0.0000420)
X3BA			1.276***	−0.000894***	−0.000690***
			(0.171)	(0.0002136)	(0.000133)
HRA			1.488***	−0.000957***	−0.000952***
			(0.0602)	(0.0000756)	(0.0000467)
Constant	0.521***	−376.685***	−473.298***	0.600***	0.517***
	(0.0218)	(26.448)	(27.010)	(0.0473)	(0.0128)
Observations	596	596	596	596	596
R-squared	0.888	0.938	0.935	0.828	0.824

Standard Errors in Parentheses

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

Team and season level fixed effects included

---

[4]See Albert and Bennett (2001) or Albert (2016) for examples using this approach.

[5]http://www.seanlahman.com/baseball-archive/statistics/

12

The first column reports the estimates for the relationship between win percentage, runs and runs against. The second column estimates runs as a function of the five main batting events produced by the batters, and the third column reports the analogous estimates for runs against. Column (4) combines these estimates to relate batting events to win percentage, while column (5) constrains the for and against coefficients for each event type to be of equal and opposite sign.

While these regressions explain approximately 94% of the variation in both runs and runs against by using our five standard batting events, we have omitted other variables (i.e. fielding, age, etc.) that are likely related to team success. We address these variables in more detail when it comes to testing our main hypothesis.

2.3.2   Player Statistics

We now want to estimate the contribution to runs of each batter based on their production of events such as walks, singles, etc., across the season, using the approach developed by Lindsey (1963). Each event type has a run value, which we can estimate based on the run expectancy matrix. The intuition behind the run expectancy matrix is that batting statistics are context dependent – the value of an action depends on the state of the half inning in which it occurs. Whenever a batter steps up to the plate, the half inning is in one of 24 potential states, which are defined by the states of the bases and the number of outs – see Figure 2.1.

Figure 2.1: The Run Expectancy Matrix

Base States	Outs		
	0	**1**	**2**
000			
100			
010			
110			
001			
101			
011			
111			

Thus base state 000 means no runners on base, 100 means a runner on first and no runners on second or third, 010 means no runner on first, a runner on second and no runner on third, and so on. The run expectancy matrix is then populated by finding the average number of runs scored from any given state to the end of the half inning, for all events in an entire season. This can be done using event data available at Retrosheet.org or in recent years from files made available by MLBAM.[6] The event level data used to derive the run expectancies and run values in this paper comes from Retrosheet.

---

[6]https://www.retrosheet.org/game.htm;
https://www.rdocumentation.org/packages/mlbgameday/versions/0.2.0

Table 2.2: The Run Expectancy Matrix for the 2011 MLB Season

	0	1	2
000	0.471	0.255	0.097
100	0.835	0.496	0.218
010	1.058	0.650	0.309
110	1.414	0.874	0.422
001	1.454	0.937	0.317
101	1.753	1.150	0.488
011	1.930	1.339	0.541
111	2.172	1.475	0.761

Table 2.2 is the run expectancy matrix for the 2011 MLB season, based on 191,864 events. Thus, for example, the average number of runs scored, from the state of no outs and no runners on base until the end of the half inning, was 0.471. The number of runs scored on average with no runners on base and two outs until the end of the half inning was 0.097. Not surprisingly, run expectancy tends to increase when there are more runners on base, when the runners are further advanced, and when there are fewer outs. There are some variations in the run expectancy matrices from year to year, but generally these variations are quite small. Over the timeframe 1996-2015, the standard deviation in run expectancy for the season was below 0.05 runs for 17 out of the 24 base states and ranged from 0.008 runs to 0.102 runs overall. As expected, the rarer, less frequent base states (e.g. bases loaded and no outs) resulted in the most variation in run expectancy from year to year.

From the run expectancy matrix we can calculate the average run expectancy of a given event (a single, a walk, a home run, an out, etc.) in a season. For any event, there is a change in run expectancy which is defined as the difference between the run expectancy immediately before the event (the start state) and run expectancy when the event ends (the end state).[7]

---

[7]Technically, to define all possible run values we need to define the run value of end states where there are three outs (since these end states occur). For each of the eight possible states (relating to the configuration of the base runners) the run expectancy is zero, since the half inning is over.

For example, for the 2011 season, for any given event, the change in run expectancy can be derived simply by taking the difference in cell values, identifying the start state and the end state in Table 2.2. The run value of an event can then be defined as the runs scored during the event, plus the difference in run expectancy at the beginning of an event and at the end of the event:

$$Run\ Value = Runs\ Scored\ During\ Event +$$

(2.4) $$Expectancy\ at\ the\ End\ of\ the\ Event -$$

$$Run\ Expectancy\ at\ the\ Beginning\ of\ the\ Event.$$

Using the empirical frequencies we can construct a table of run values for a given event, depending on the starting states. This is illustrated for the home run event in the 2011 season in Table 2.3.

Table 2.3: Run Value of Home Runs in 2011

	0	1	2
000	1.000	1.000	1.000
100	1.636	1.759	1.879
010	1.413	1.604	1.788
110	2.057	2.381	2.675
001	1.017	1.317	1.780
101	1.718	2.105	2.609
011	1.541	1.916	2.556
111	2.299	2.780	3.336

The run value of a home run, with no one on base at the beginning of the event, is always 1, which is just the run scored, given that there is no one on base at the end of the event (i.e. the bases states are unchanged). At the other extreme, the run value of bases loaded and two outs (111 2) is 3.336 runs, which is the value of the four runs scored plus the change

16

in value of the base states.[8]

We can also calculate the average value of an event across all starting states. Thus, for example, the run value of a home run in a given season is the weighted sum of the run values of a home run across starting states, where the weights are based on the relative frequencies of each starting state for that event. Based on run values in Table 2.3 and the relative frequency of home runs from each starting state in 2011, the average run value of a home run was 1.39 in that season.

Table 2.4 shows the weighted run values of our main batting outcomes for each of the seasons in our data. These values are often referred to as "linear weights" in the literature.[9] It is clear from Table 2.4 that these values are very stable over time. This is not the same as saying that the game never changes. Over this period there have clearly been many innovations in baseball, not least those associated with *Moneyball* and the increasing use of data analytics. The first half of the data also roughly coincides with the "steroid era", which increased batting power and thus the production of home runs. While the aggregate production of particular events may have changed over time, Table 2.4 indicates that the frequency with which particular events are converted into runs has remained very stable over this period.

---

[8]At first glance, this might seem counterintuitive- four runs were scored, but the run value is less than four. The intuitive explanation is that while four runs were scored, the event itself should not be given the full credit for these runs, since hitting the home run was not responsible for getting the three runners on base at the beginning of the event – that achievement should be credited to earlier events. Note that at the beginning of the event, when bases were loaded, the run expectancy was very high, but at the end of the event there are no runners on base, so that run expectancy has fallen considerably.

[9]Albert and Bennett (2001) attribute this coinage to Peter Palmer (see p. 200).

Table 2.4: Run Value Weights by Season for Batting Events

Season	Walk	Single	Double	Triple	Home Run	Out
1996	0.331	0.485	0.784	1.105	1.403	-0.302
1997	0.307	0.465	0.761	1.083	1.393	-0.284
1998	0.312	0.469	0.780	1.014	1.400	-0.285
1999	0.311	0.477	0.789	1.059	1.408	-0.302
2000	0.332	0.482	0.765	1.085	1.406	-0.307
2001	0.298	0.460	0.778	1.084	1.380	-0.283
2002	0.303	0.466	0.755	1.052	1.398	-0.279
2003	0.307	0.466	0.775	1.080	1.391	-0.284
2004	0.307	0.462	0.786	1.041	1.396	-0.287
2005	0.295	0.458	0.768	1.056	1.412	-0.277
2006	0.317	0.467	0.766	1.070	1.389	-0.290
2007	0.310	0.468	0.798	1.044	1.406	-0.289
2008	0.312	0.460	0.772	1.081	1.405	-0.281
2009	0.304	0.459	0.762	1.004	1.392	-0.278
2010	0.299	0.451	0.763	1.076	1.404	-0.266
2011	0.289	0.442	0.736	1.064	1.392	-0.255
2012	0.284	0.441	0.747	1.039	1.396	-0.257
2013	0.285	0.439	0.740	1.035	1.371	-0.250
2014	0.283	0.437	0.739	1.054	1.400	-0.245
2015	0.303	0.442	0.743	1.031	1.386	-0.257
mean	0.304	0.459	0.764	1.056	1.396	-0.277
standard deviation	0.014	0.014	0.018	0.027	0.010	0.018

To confirm that our calculated run value weights accurately represent the contribution of each outcome to runs scored, we run an OLS regression of team runs scored on the team totals for each outcome weighted by the appropriate run value for the season from Table 2.4 for the seasons 1996-2015. For example, if a team had a total of 500 walks in 1996, the team's run value contribution of walks for the season is estimated by multiplying the run value weight for walks in 1996 (0.331) by 500. We calculate the run values for each outcome that contributes to the run scoring process (walks, singles, doubles, triples and home runs) at the team level for every season for 1996-2015. We also run a similar OLS regression model replacing runs scored and our aggregate run values with runs scored against and aggregate run values against respectively. Finally, we run both versions of this OLS model for all seasons and for the post-*Moneyball* era (post-2003) as well. In all regressions, team and

season fixed effects are included to account for potential ballpark effects and changes to the game over our timeframe. The results are shown in Table 2.5.

Table 2.5: Regressions of Runs on Run Values

Dependent Variable Seasons	Runs Scored All	Runs Scored Post-2003	Runs Scored Against All	Runs Scored Against Post-2003
RV Walk	1.127***	1.015***		
	(0.0630)	(0.0877)		
RV Single	1.086***	1.008***		
	(0.0425)	(0.0532)		
RV Double	0.950***	0.979***		
	(0.0601)	(0.0802)		
RV Triple	1.319***	1.567***		
	(0.129)	(0.162)		
RV HR	1.001***	0.995***		
	(0.0290)	(0.0384)		
RV Against Walk			1.223***	1.328***
			(0.0699)	(0.0963)
RV Against Single			1.155***	1.200***
			(0.0500)	(0.0694)
RV Against Double			0.975***	0.908***
			(0.0596)	(0.0763)
RV Against Triple			1.205***	1.516***
			(0.163)	(0.204)
RV Against HR			1.059***	1.090***
			(0.0433)	(0.0615)
Constant	−428.802***	−325.303***	−524.341***	−523.412***
	(27.520)	(40.536)	(28.270)	(41.402)
R^2	0.939	0.937	0.934	0.926
Num. obs.	596	360	596	360
F test	2.492 (p 0.0837)	0.0389 (p 0.962)	1.946 (p 0.144)	2.030 (p 0.133)

Standard Errors in Parentheses

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

F test for equality of coefficients on walks, singles, and home runs

Team and season level fixed effects included

From the results, we observe an R-squared of approximately 0.93-0.94 in all regressions, indicating a very strong relationship between runs scored and our weighted run value outcomes. We also observe that the coefficient sizes for several outcomes tend to be similar as well. In particular, the coefficients for walks, singles, and home runs look to be comparable in each regression. The coefficients for doubles appears to be consistently smaller

than all other coefficients while the coefficients for triples tend to be larger. Doubles and triples also tend to have lower levels of significance than walks, singles and home runs in these regressions. This is likely related to the fact that doubles and triples tend to be more random (and rare in the case of triples) compared to walks, singles, and home runs. For this reason, doubles and triples are generally not thought of as outcomes representing a set of skills entirely distinguishable from those captured by either singles or home runs. Given the similar coefficients for walks, singles, and home runs, we conduct an F-test for coefficient equality for these three outcomes and find that we cannot reject coefficient equality at the .05 level in any of our regressions. This result tells us that the run values for each of these three outcomes contribute similarly to team runs. In other words, the run values generated via walks and singles are similar in value to the run values generated via home runs from strictly a run scoring perspective, even though, as we see in Table 2.4, more walks and singles are required to achieve an equivalent run value to home runs.

Table 2.6 compares the average values in Table 2.4 to the estimates of our regressions in Table 2.1. Not surprisingly, most regression estimates are similar in magnitude to the values derived from the run expectancy matrix, with triples being an exception due to the limited number of triples that occur in a season (the average player in our data hit only two triples per season and only 0.5% of all plate appearances resulted in a triple). One might also note that the regression coefficients do tend to be slightly inflated compared to their corresponding run value weights. This is likely due to potential omitted variable bias present in the regression estimates, but not the structural run value estimates as explained in Turocy (2005) and Berri and Bradbury (2010). However, as Albert and Bennett (2001) point out, the two types of approaches tend to ultimately produce similar results.

Table 2.6: Run Weights from Regressions (Table 2.1) and from Run Values (Table 2.4)

	Team Runs (Table 2.1)	Team Runs Against (Table 2.1)	Run Value (Table 2.4)
Walk	0.347	0.378	0.304
Single	0.498	0.534	0.459
Double	0.730	0.739	0.764
Triple	1.402	1.276	1.056
Home Run	1.398	1.488	1.396

Given that we have a stable estimator of the run value contribution of events, we can extend this to estimate the run value contribution of each batter in each season. The run value of a batter is defined as the runs contributed by that batter in a season based on the events he produced (i.e. singles, walks, home runs, outs, etc). Following Equation (2.4), the run value of each batter event consists of the runs scored during the event plus the change in run expectancy (based on the run expectancy matrix). There are two ways to calculate the run values over a season for a given batter. We can (a) calculate the average value of an event (a walk, a single, etc) across all events and then multiply the batter's season totals in each category by the average event value, or (b) we can simply calculate sum of the values of the batter's events in each category across the season. As a result, the weights applied to each respective event type will be identical for all players in a season (as given in Table 2.4) under approach (a). On the other hand, approach (b) is equivalent to calculating the average run value of each event type separately for every player during the season, and then applying these player-specific weights to the player's event totals for the season.

The first method, which we will refer to as the "season" weights approach, tends to smooth out the player's contributions, and thus removes some of the random factors that influence an individual event which are not under the control of the batter (e.g. good/bad fielding or base running). It also allows us to normalize for any differences that may arise between players due to their batting position in the lineup. As a check, we calculated the

average run value for each batting event type included in Equation (2.2) by batting lineup position and conducted an ANOVA test and subsequent Tukey HSD tests to determine if significant differences exist between average run values by batting lineup position. We only find modest differences in average run value by batting event type across all batting lineup positions. Out of the 36 pairwise combinations, the average run values for home runs had the highest number of significant differences with seven, as the leadoff position had a significantly lower average run value than all other lineup positions. Equivalent tests for walks, singles and doubles produced six, five, and five pairwise significant differences respectively, while these tests did not produce any significant differences for triples. We conclude that run value differences between batting lineup position are not expected to impact on our results, but that approach (a) will nevertheless allow us to normalize for some of these differences.

Additionally, while it is a maintained assumption that batters are not able to produce events tailored to the context of the game, one could think of the difference between the two as a measure of clutch performance for the player during the season. If the player's run value calculated using approach (b), which we refer to as the "player" weights approach, is higher than his run value calculated using approach (a), this means that the player accrued his batting outcomes in more valuable run-scoring situations than the average player. However, given that the best batters will have over 100 events in a season, these alternatives are likely to generate very similar rankings.

We can use the run value of batter contributions to estimate a salary regression, analogous to the salary regressions in Table 3 of Hakes and Sauer (2006). We hypothesize that in an efficient market, player run value contributions should be valued equally, regardless of the type of event. Thus, for example, a run contributed by a batter as a result of walks should be valued equally with a run contributed by means of a home run. If we take the batter event statistics for a season, and then weight them using the run value estimates for the season described above, we have an appropriately weighted estimate of the run contribution

for each batter from each event type. For example, from Table 2.4 we observe that three doubles have an approximately equal run value (1.377) compared to a home run (1.396). Thus two batters who differ only because (a) one produced three more doubles and (b) the other produced one more home run, should be valued equally.

To conduct our test of market efficiency, we begin by proposing the following basic pooled OLS regression model:

$$(2.5) \qquad RelSal \ = \ \beta_0 \ + \ \beta_1 RV_{Walk} \ + \ \beta_2 RV_{Single} \ + \ \beta_3 RV_{XBH} \ + \ \beta_4 RV_{HR} \ + \epsilon,$$

where our dependent is the natural log of player salary indexed to the average by season, and our independent variables are each run values for the season aggregated by outcome. We use a panel dataset for our salary regressions where our salary data is obtained at the player and season level from the Lahman Database, and our run value outcome variables are aggregated at the player and season level as well and derived using the event level data from Retrosheet. We index salaries to the average by season to account for any changes to the salary distribution that occurred over our timeframe.[10] Additionally, we regress current season salary on the previous season's batting statistics, given that the most reliable guide to future performance is likely to be past performance. We combine a player's run values accumulated via doubles and triples together in one variable (RV_{XBH}), since triples are very rare and both outcomes tend to be more random in nature compared to walks, singles and home runs (the year-to-year correlation for extra-base hits for all players in our data is approximately 0.48 compared to 0.70, 0.57, and 0.70 for walks, singles, and home runs respectively). This regression is run on several subsets of our data to determine how the

---

[10]In particular, indexing player salaries to the average by season will help to normalize salaries under different collective bargaining agreements (CBAs) that were reached over our timeframe. The seasons in which new CBAs went into effect include 1997, 2003, 2007, and 2012.

labor market valued these outcomes in the years surrounding the publication of *Moneyball*. First, we run the regression for all players over the timeframe 1997-2016, excluding salaries in the season 2004. Given that we regress salary on the previous season's batting statistics, we omit the observation for 2004, since this confounds salary paid after the publication of *Moneyball* with statistics produced before the publication of *Moneyball*.

Next, we subset the data and rerun the regression for free agent players only over this timeframe to avoid any potential inefficiencies that may exist in the "all players" regression due to the limited bargaining power of reserve clause and arbitration eligible players. We define a free agent to be a player that meets at least one of the following two conditions: (i) The player has at least seven years of experience, or (ii) the player has been signed as a free agent in the offseason at some point prior to the player's "salary season". In this classification, we are using years of experience in the league as a proxy for MLB service time, based on the year in which they made their MLB debut. Under the current rules, players are typically eligible for free agency the season after which they reach their sixth year of MLB service time. We decide to use seven years of experience (rather than six) as our free agent classification, to avoid including players with six years of experience but less than six years of service time. This scenario could happen if the player made his debut in middle or end of a season, or if the player gets sent down to the minor leagues for an extended period after making his MLB debut. To make sure we capture players who truly were eligible for free agency after their sixth year of experience, we include all players who meet condition (ii) in our free agent classification as well. Player debut data was obtained from the Lahman Database while offseason signing data was obtained from Retrosheet's transaction data.[11]

All regressions will be run using the two approaches to calculate run values for players: (a) the "season weights" approach and (b) the "player weights" approach. We follow Hakes and Sauer (2006, 2007) in restricting our analysis to batters with at least 130 at bats in

---

[11]https://retrosheet.org/transactions/index.html

a season. We also follow the approach of Holmes et al (2018) and cluster standard errors by player when running this regression. This is done since our data includes repeat player observations for the seasons over our timeframe. To test our hypothesis of equal returns to the weighted contribution of batting events, we use an F-test for the equality of coefficients for walks, singles, and home runs. In an efficient market, we would not expect to reject the equality of these coefficients in our salary regression, given that we did not reject our F-test for these coefficients in our regressions on Table 2.5.

The first two columns of Table 2.7 refer to all players in our sample, while the final two columns focus on free agents.

Table 2.7: Salary Regressions for the Entire Sample Period: 1997-2016 (excluding 2004)

Dependent Variable	RelSal	RelSal	RelSal	RelSal
Season or Player Weights	Season	Player	Season	Player
Player Type	All	All	Free Agent	Free Agent
Run Value Walk	0.00239***	0.00220***	0.00125***	0.00122***
	(0.000213)	(0.000205)	(0.000219)	(0.000214)
Run Value Single	0.00174***	0.00161***	0.00141***	0.00135***
	(0.000113)	(0.000107)	(0.000120)	(0.000112)
Run Value XBH	−0.000575***	−0.000145	0.000945***	0.00102***
	(0.000181)	(0.000163)	(0.000186)	(0.000166)
Run Value HR	0.00177***	0.00167***	0.00162***	0.00152***
	(0.000100)	(0.0000921)	(0.000101)	(0.0000909)
Constant	0.892***	0.892***	0.931***	0.933***
	(0.00240)	(0.00239)	(0.00346)	(0.00343)
Observations	6488	6488	2921	2921
R-squared	0.370	0.366	0.498	0.492
F test	2.94 (p 0.055)	2.70 (p 0.069)	1.30 (p 0.275)	0.86 (p 0.425)

Standard Errors Clustered by Player in Parentheses

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

F test for equality of coefficients on walks, singles, and home runs

From Table 2.7, when we include all players (with more than 130 at bats) we find that the hypothesis of coefficient equality can only be rejected at the .10 significance level, and when we restrict the data to free agents, the hypothesis cannot be rejected. This is a striking

result given the conclusions drawn from previous research. For free agents, it suggests that earnings accurately reflected their runs contribution to the team throughout period 1997-2016 (excluding 2004), including the period before the publication of *Moneyball*. The weaker efficiency we observe for all players is not so surprising. Previous research has suggested that reserve clause players (less than three years of service time) and arbitration eligible players (between three to six years of service time) earn lower salaries than free agent players (at least six years of service time), allegedly enabling teams to recoup development costs (e.g. Miller (2000)). Hakes and Sauer (2006) identified significant dummy variables related to contract status, while Brown et al (2017) found very different returns to batting performance contingent on contract status. Free agent salaries are more likely to reflect the market returns to specific skills.

We now consider whether there were any changes to these relationships before and after the publication of *Moneyball* in 2003. We estimate the same regressions using the data for 1997-2003 and then for 2005-2016.[12] The results are reported in Table 2.8. For this analysis we focus only on free agents with more than 130 at bats in a season.

---

[12]Given that we regress salary on the previous season's batting statistics, we omit the observation for 2004, since this confounds salary paid after the publication of Moneyball with statistics produced before the publication of *Moneyball*.

Table 2.8: Salary Regressions Dividing the Free Agent Sample into Pre-*Moneyball* (1997-2003) and Post-*Moneyball* Periods (2005-2016)

Dependent Variable	RelSal	RelSal	RelSal	RelSal
Pre or Post *Moneyball*	Pre	Pre	Post	Post
Season or Player Weights	Season	Player	Season	Player
Player Type	Free Agent	Free Agent	Free Agent	Free Agent
Run Value Walk	0.00113***	0.00114***	0.00179***	0.00172***
	(0.000292)	(0.000282)	(0.000283)	(0.000278)
Run Value Single	0.00154***	0.00145***	0.00142***	0.00138***
	(0.000170)	(0.000153)	(0.000146)	(0.000139)
Run Value XBH	0.000718*	0.000785**	0.000785**	0.000892***
	(0.000272)	(0.000243)	(0.000238)	(0.000213)
Run Value Home Run	0.00168***	0.00158***	0.00156***	0.00148***
	(0.000126)	(0.000113)	(0.000143)	(0.000132)
Constant	0.920***	0.923***	0.935***	0.936***
	(0.00518)	(0.00513)	(0.00433)	(0.00432)
Observations	1150	1150	1771	1771
R-squared	0.553	0.546	0.482	0.477
F test	0.975 (p 0.381)	0.749 (p 0.476)	0.885 (p 0.415)	0.683 (p 0.507)

Standard Errors Clustered by Player in Parentheses

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

F test for equality of coefficients on walks, singles, and home runs

Not surprisingly, these results are closely comparable to the free agent results in Table 2.7. The returns to the run value for each type of event are relatively stable, although some variations are apparent. The hypothesis of coefficient equality for walks, singles, and home runs cannot be rejected in any regression when looking at the pre and post-*Moneyball* periods separately. It is also notable from the first two rows of Table 2.8 that the returns to the run value of a walk was lower than the returns to the run value of a single and the run value of a home run in the pre-*Moneyball* era, but was higher in the post-*Moneyball* era. While these results are consistent with the *Moneyball* narrative, they do not yield statistically significant differences in our test for market efficiency.

One key difference between the regressions reported so far and previous studies is that, in our formulation, we have not included any variables other than those directly related

to batting performance. Hakes and Sauer (2006, 2007) control for the number of plate appearances and include dummies for player fielding positions and contract status. Brown et al (2017) include an experience variable, Holmes et al (2018) include plate appearances and age, and Duquette et al (2019) include measures of fielding, pitching and base running performance. Omission of confounding variables leads to omitted variable bias, so that the estimated coefficients are biased. However, adding confounders which are correlated with the existing regressors could lead to overfitting.

Table 2.9: Correlation Coefficients for Run Value Events, Plate Appearances and Age

	Salary	Run value walk	Run value single	Run value xbh	Run value home run	Plate appearances	Age
Salary	1						
Run value walk	0.404	1					
Run value single	0.379	0.504	1				
Run value xbh	0.414	0.559	0.717	1			
Run value home run	0.486	0.633	0.383	0.533	1		
Plate appearances	0.498	0.707	0.875	0.821	0.625	1	
Age	-0.055	0.021	-0.050	-0.098	-0.080	-0.089	1

Table 2.9 shows the correlation coefficients for the regressors in Equation (2.5) (free agents, all seasons 1996-2015) with plate appearances and age added. Both age and plate appearance data are also obtained from the Lahman Database. It is clear that plate appearances are very highly correlated with our run value variables, where the correlations are higher than the correlation with salary, the dependent variable. Age, on the other hand, does not appear to be correlated with any of the explanatory variables. We therefore add age to our regression model but omit plate appearances due to its high correlation with run values.

The high correlations observed between player plate appearances and run values across batting event variables are expected, since players that are given the most playing time are the ones that will tend to accumulate higher run values. Therefore, despite the exclusion

of total plate appearances as a variable in our regression model, plate appearances can be thought of as "embedded" in our run value variables due to the strong link between plate appearances and run values. Logically, the value of a player to his team is his run value contribution, rather than his rate of run value contributions. Each plate appearance has an opportunity cost, and the standard run value contribution of a replacement player, a player who is outside of the 750 active major league players since each the 30 teams are allowed 25 player active on their roster (Baumer et al, 2015), is negligible. Woolner (2007) provides an example, comparing an average player who appears in every game to a star player who appears in only one sixth of all games and whose place has to be taken by a replacement for the rest of season, showing that the average player is actually worth more to the team, in terms of runs, than the star + replacement.

Hence, rather than include total plate appearances in our model, we include a measure of the season to season variation in a player's past plate appearances to account for the risk that a team incurs by signing a player with recent playing time volatility (this could be due to injury, suspensions, etc.). For this variable we calculate the coefficient of variation for a player's plate appearances measured over the previous three seasons. Therefore, we run the following OLS regression model for all seasons and for the pre-*Moneyball* and post-*Moneyball* eras separately in a more robust version of our original model:

(2.6)
$$RelSal = \beta_0 + \beta_1 RV_{Walk} + \beta_2 RV_{Single} + \beta_3 RV_{XBH} + \beta_4 RV_{HR} + \beta_5 CV_{PA} + $$
$$\beta_6 Age + \beta_7 Age^2 + \sum_i \alpha_i POS_i + \sum_i \gamma_j Team_j + \epsilon,$$

where our coefficient of variation measure for plate appearances, age, age-squared, player position fixed effects (to capture the different levels defensive ability required for different positions), and team fixed effects (to capture the impact of factors such as market size and

29

location on salary) are added. Additionally, we follow the approach of Holmes et al (2018) again and cluster standard errors by player when running this regression. Ultimately this regression model is similar in structure to the models from previous research pertaining to *Moneyball* with several elements of previous research integrated into this model. Of course, the main difference with our approach is the usage of our run value variables as player performance indicators as opposed to OBP and SLG or metrics that serve as proxy variables for OBP and SLG. Table 2.10 reports the OLS regression results for Equation (2.6) for all seasons (1997-2016, excluding 2004), the pre-*Moneyball* period (1997-2003), and the post-*Moneyball* period (2005-2016) separately.

Table 2.10: Selected Coefficients - Adding Variables Other than Run Values

Dependent Variable	RelSal	RelSal	RelSal	RelSal	RelSal	RelSal
Seasons	All	All	pre-MB	pre-MB	post-MB	post-MB
Season or Player Weights	Season	Player	Season	Player	Season	Player
Player Type	Free	Free	Free	Free	Free	Free
	Agent	Agent	Agent	Agent	Agent	Agent
Run Value Walk	0.00121***	0.00121***	0.00115***	0.00121***	0.00177***	0.00171***
	(0.000209)	(0.000210)	(0.000255)	(0.000250)	(0.000282)	(0.000286)
Run Value Single	0.00124***	0.00115***	0.00140***	0.00127***	0.00119***	0.00112***
	(0.000117)	(0.000110)	(0.000182)	(0.000168)	(0.000148)	(0.000141)
Run Value XBH	0.000936***	0.000998***	0.000775**	0.000830***	0.000743**	0.000832***
	(0.000178)	(0.000159)	(0.000270)	(0.000244)	(0.000232)	(0.000211)
Run Value Home Run	0.00152***	0.00144***	0.00157***	0.00149***	0.00140***	0.00133***
	(0.0000989)	(0.0000897)	(0.000118)	(0.000108)	(0.000144)	(0.000135)
PA CV	−0.0271***	−0.0292***	−0.0208***	−0.0230***	−0.0274***	−0.0294***
	(0.00454)	(0.00453)	(0.00573)	(0.00563)	(0.00654)	(0.00658)
Age	0.0354***	0.0338***	0.0435***	0.0419***	0.0271***	0.0253***
	(0.00657)	(0.00631)	(0.0105)	(0.0105)	(0.00662)	(0.00622)
Age2	−0.000542***	−0.000520***	−0.000668***	−0.000644***	−0.000421***	−0.000395***
	(0.000101)	(0.0000968)	(0.000164)	(0.000164)	(0.0000994)	(0.0000929)
Fielding: Second Base	−0.00941	−0.00567	−0.0140	−0.00809	−0.00398	−0.00183
	(0.00559)	(0.00566)	(0.00725)	(0.00753)	(0.00739)	(0.00741)
Fielding: Third Base	0.00389	0.00467	−0.00141	−0.0000720	0.00854	0.00876
	(0.00566)	(0.00567)	(0.00775)	(0.00768)	(0.00716)	(0.00708)
Fielding: Catcher	0.00224	0.00293	0.00607	0.00763	0.000989	0.000958
	(0.00519)	(0.00521)	(0.00768)	(0.00777)	(0.00653)	(0.00646)
Fielding: Designated Hitter	−0.00413	−0.00454	−0.00586	−0.00523	−0.00223	−0.00330
	(0.00608)	(0.00611)	(0.00841)	(0.00853)	(0.00824)	(0.00819)
Fielding: Outfield	0.00310	0.00514	−0.00132	0.00111	0.00676	0.00806
	(0.00418)	(0.00422)	(0.00541)	(0.00542)	(0.00561)	(0.00561)
Fielding: Shortstop	0.00608	0.00983	0.00484	0.00919	0.00703	0.00995
	(0.00579)	(0.00597)	(0.00897)	(0.00929)	(0.00732)	(0.00744)
Constant	0.367***	0.396***	0.244	0.272	0.500***	0.534***
	(0.106)	(0.102)	(0.167)	(0.167)	(0.110)	(0.104)
Observations	2921	2921	1150	1150	1771	1771
R-squared	0.536	0.529	0.598	0.592	0.533	0.526
F test	1.94	2.05	0.988	0.903	1.87	2.00
F test p-value	0.147	0.132	0.376	0.408	0.158	0.139

Standard Errors Clustered by Player in Parentheses
$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$
F test for equality of coefficients on walks, singles, and home runs
Team level fixed effects included

The results in Table 2.10 are entirely consistent with Tables 2.7 and 2.8. The estimates for the run value coefficients are similar in magnitude, and we fail to reject our F-test for coefficient equality in every regression once again. We observe that our measure of plate appearance variation is negative and significant across all models. This intuitively makes sense as higher levels of plate appearance variation reflect higher levels of risk and uncertainty

associated with the signed player. Thus, teams are more likely to apply a discount to such a player. Age and age squared have the right signs, and we do not find significant fixed effects for any player position. Overall, we take these results as confirmation that our main results are robust to the inclusion of additional covariates that are not highly correlated with the run value variables.

Thus far we have excluded extra-base hits from our test for coefficient equality due to the year-to-year randomness of extra-base hits and the difficulty in separating the skills associated with extra-base hits from the skills associated with our other batting outcomes. One potential concern in comparing our results to the results of previous research using OBP and SLG is that both OBP and SLG include extra-base hits in their formulation. To address this concern, we run an alternative specification of Equation (2.6) where we aggregate a player's run values for walks and singles into a metric we call "RV On-Base", and we aggregate a player's run values for doubles and home runs into a metric we call "RV Power". These two metrics will replace the four disaggregated run value independent variables from our original specification in Equation (2.6).

Recall that the formula for OBP weights all batting outcomes the same in its formulation. Based on the run value weights for batting outcomes that we calculated in Table 2.4, this means that walks and singles are weighted the most relative to their actual run values in OBP. Thus, we combine the run values of walks and singles into our "RV On-Base" metric to serve as a run value version of OBP. Likewise, recall that SLG places more value on doubles, triples, and home runs in its formulation. Therefore, we combine the run values of doubles and home runs into our "RV Power" metric to serve as a run value version of SLG. Note that we continue to exclude triples from this metric since triples are rare and are not thought of as being highly correlated with a specific set of skills. In this alternative specification we then complete our F-test for coefficient equality on the coefficients of RV On-Base and RV Power to determine whether equal returns to run values continue to be observed for these

"run value versions" of OBP and SLG. The results are included in Table 2.11.

Table 2.11: Selected Coefficients - Power and On-Base Run Values

Dependent Variable	RelSal	RelSal	RelSal	RelSal	RelSal	RelSal
Seasons	All	All	pre-MB	pre-MB	post-MB	post-MB
Season or Player Weights	Season	Player	Season	Player	Season	Player
Player Type	Free	Free	Free	Free	Free	Free
	Agent	Agent	Agent	Agent	Agent	Agent
Run Value On-Base	0.00118***	0.00109***	0.00124***	0.00114***	0.00123***	0.00115***
	(0.0000836)	(0.0000844)	(0.000111)	(0.000113)	(0.000109)	(0.000112)
Run Value Power	0.00140***	0.00134***	0.00141***	0.00136***	0.00131***	0.00126***
	(0.0000720)	(0.0000714)	(0.0000905)	(0.0000900)	(0.0000990)	(0.0000972)
PA CV	−0.0261***	−0.0291***	−0.0201***	−0.0230***	−0.0259***	−0.0287***
	(0.00456)	(0.00456)	(0.00572)	(0.00563)	(0.00663)	(0.00668)
Age	0.0346***	0.0339***	0.0428***	0.0413***	0.0264***	0.0259***
	(0.00649)	(0.00624)	(0.0106)	(0.0104)	(0.00646)	(0.00621)
Age2	−0.000530***	−0.000520***	−0.000658***	−0.000635***	−0.000409***	−0.000402***
	(0.0000998)	(0.0000957)	(0.000165)	(0.000162)	(0.0000967)	(0.0000925)
Fielding: Second Base	−0.00896	−0.00714	−0.0135	−0.0100	−0.00571	−0.00458
	(0.00553)	(0.00564)	(0.00712)	(0.00742)	(0.00744)	(0.00746)
Fielding: Third Base	0.00398	0.00405	−0.000411	0.0000234	0.00674	0.00674
	(0.00561)	(0.00566)	(0.00781)	(0.00775)	(0.00715)	(0.00711)
Fielding: Catcher	0.00235	0.00246	0.00715	0.00785	−0.000965	−0.00108
	(0.00513)	(0.00519)	(0.00765)	(0.00776)	(0.00653)	(0.00650)
Fielding: Designated Hitter	−0.00384	−0.00430	−0.00592	−0.00529	−0.00182	−0.00298
	(0.00608)	(0.00613)	(0.00844)	(0.00852)	(0.00822)	(0.00820)
Fielding: Outfield	0.00435	0.00457	0.000257	0.00104	0.00639	0.00637
	(0.00415)	(0.00424)	(0.00542)	(0.00548)	(0.00567)	(0.00569)
Fielding: Shortstop	0.00674	0.00830	0.00655	0.00802	0.00426	0.00610
	(0.00577)	(0.00595)	(0.00885)	(0.00922)	(0.00740)	(0.00746)
Constant	0.379***	0.394***	0.253	0.279	0.511***	0.524***
	(0.105)	(0.101)	(0.167)	(0.166)	(0.108)	(0.104)
Observations	2921	2921	1150	1150	1771	1771
R-squared	0.537	0.528	0.599	0.590	0.530	0.523
F test	2.55	3.19	0.952	1.39	0.214	0.346
F test p-value	0.112	0.076	0.321	0.242	0.644	0.558

Standard Errors Clustered by Player in Parentheses
$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$
F test for equality of coefficients on on-base and power
Team level fixed effects included

Once again, the findings from our previous model are confirmed when we replace our disaggregated run value metrics with our RV On-Base and RV Power metrics. Our F-test remains insignificant at the .05 level for all of our regressions and only one regression is significant at the .10 level (player weights, all seasons). Therefore, our previous conclusion of teams paying players according to their run value contributions holds when using this broader categorization of outcomes more compatible with OBP and SLG.

To further check the robustness of our model, we ran several alternate specifications of Equation (2.6). We find that all F-test results for coefficient equality are robust to defining free agents using six years of experience rather than seven, replacing age and age-squared with experience and experience-squared in Equation (2.6), including the salary season 2004 in our data, and using z-score of the natural log of salary (standardized by season) as our dependent variable. Additionally, we also ran an alternative specification including player fixed effects as well in Equation (2.6) to account for additional factors such as star quality, fielding ability and multi-year contract signings. We continue to find that all F-test results for coefficient equality are insignificant for the models analogous to those run in Table 2.10. Overall, the results for these alternative specifications support our finding of equal returns to event types in the labor market over our entire timeframe and confirm the robustness of these results.

2.4 Conclusions

The publication of *Moneyball* in 2003 and subsequent research has questioned the efficiency of the baseball labor market, suggesting that prior to 2003 the market undervalued certain characteristics such as the capacity of batters to draw walks. In this paper we have advanced a test of the efficiency hypothesis using the concept of the run expectancy matrix and the run value, a measure of the contribution of different batting events (walks, singles, home runs, etc) to generating runs and therefore wins. While much of the focus on testing for market efficiency in the baseball labor market has focused on OBP and SLG, we argue that run values provide us with a more rigorous approach to measuring player productivity and allow us to determine a player's contributions to team output more explicitly. This is one of the primary reasons why the sabermetrics community has shifted towards a run value based approach in many advanced performance metrics, including the popular wins above

replacement (WAR) metric (Baumer et al, 2015). While OBP and SLG are useful as quick and easy calculations that are correlated with player ability, the run value approach takes the measurement of player productivity one step further and measures a player's contributions to runs based on a more granular level of a player's performance outcomes. This paper bridges the gap that exists between the sabermetrics community and the current sport economics literature by utilizing this run value approach to test for labor market efficiency in Major League Baseball.

Our results are striking. Omitting the category of extra-base hits, which are more random events in nature, we find that we cannot reject the hypothesis of equal returns to event types in a salary regression for free agents over the entire period 1997-2016 (excluding 2004). When we divide the data into pre- and post-*Moneyball* (publication) eras, we find some evidence of an increase in the coefficient for walks in the post-*Moneyball* era, but this increase does not yield statistically significant differences in the returns to the run values of walks, singles, and home runs. To the extent that we observe increased efficiency post-*Moneyball*, the difference seems marginal. These results are also robust to the inclusion of covariates relating to age, variation in plate appearances and fielding position. The result of equal returns to run values also holds when replacing our disaggregated run value independent variables with "run value versions" of OBP and SLG, where we aggregate the run values of player walks and singles in the metric RV On-Base and the run values of player doubles and home runs in the metric RV Power.

There are a number of interesting avenues for future research. We have not here analyzed the contribution of pitchers, but this is clearly feasible since each event is associated with a pitcher, just as it is with a batter. The data sources would also enable us to attribute run values to fielding and base running. Another potentially interesting application of the data could be in measuring the efficiency of player trades.

Chapter III

On the Efficiency of Trading Intangible Fixed Assets in Major League Baseball

3.1 Introduction

Market efficiency is usually framed in terms of exploitation of gains from trade in economic analysis. Testing the efficiency of labor markets is beset by two problems. First, whilst labor is the most important input in the production process, it is typically not a marketable asset that can be bought and sold, since our laws prohibit indenture and slavery. Second, to the extent that we can study the efficiency of labor market contracts between employers and workers, most labor activity takes place behind a veil of incomplete information, which makes labor productivity itself hard to assess, either by the employer or by economic researchers. Moreover, incomplete information rationalizes incentive contracts which do not guarantee ex-post efficiency, but which are ex-ante efficient on average.

Professional baseball in the US allows us in theory to overcome both of these obstacles. First, since 1946 the IRS has permitted baseball teams to treat player contracts as intangible fixed assets, and these assets are regularly traded among the teams. Second, worker productivity is observable at regular intervals and quantifiable with a high degree of precision. We

focus on batters, and we develop a novel measure of player value which is based on the mean and variance of game level production by batters, where production is measured using the outcome of each player's plate appearances during a game. Outcomes that advance players around the bases (i.e. walks, singles, home runs, etc.) contribute to the team's production of runs (by which teams win games).

We test the efficiency of a set of player trades where information on productivity and the circumstances surrounding the trade is well defined. Our data covers the period 1994-2016 and only includes players traded in mid-season. For this subset we find that 47% of trades were efficient by our measure until 2010, while since then around 58% of trades have been efficient. Moreover, since our measure of efficiency is subject to random shocks, we show in a Monte-Carlo simulation that on average, a truly efficient trade is expected to be counted as efficient by our measure approximately 64% of the time.

Thus it appears that player trading in baseball has become more efficient in recent years. We attribute this to the dramatic expansion of statistical analysis in baseball, often referred to as sabermetrics, and data analytics in sports in general. A key moment in the development of baseball was the book *Moneyball*, published in 2003, which attributed the success of the Oakland Athletics under Billy Beane to the use of statistical analysis in the process of player evaluation. Prior to this, few teams had a data analytics department, while now all teams do. While the application of advanced statistical analysis in baseball has been around since the 1960s, advances in data collection and computing power have enabled analysts to estimate performance models that were not previously possible.

3.2   Literature Review

A great deal of research in finance has been devoted to financial portfolio performance and efficiency. Much of this research builds upon the portfolio theory work of Markowitz (1952).

In this seminal paper, a framework is developed for which an investor can maximize the mean return of a portfolio while minimizing its variance through diversification. Tobin (1958) builds upon this framework by showing that an optimal investment choice can be identified for any level of risk preference by investing in a combination of the efficient portfolio and the risk-free asset. Markowitz (1959) then goes on to extend his previous work by incorporating an investor's utility function with risk preferences into the original framework. Sharpe (1964) showed that the results derived in previous research are consistent with what one would expect under the market equilibrium of asset pricing. Sharpe (1966) then formalizes the relationship between portfolio mean and variance by defining a portfolio by its ratio of expected return to risk (variability), where an efficient portfolio is one where this ratio is maximized.

The application of these seminal works in portfolio theory is extensive in the finance literature. In particular, the mean-variance tradeoff has been a focus of several studies. Best and Grauer (1991) investigate the sensitivity of efficient portfolios to changes in asset means. Fletcher (2009) conducts an empirical analysis of portfolio risk reduction using the mean-variance framework. Chiu and Wong (2011) apply the mean-variance framework to the portfolio selection of cointegrated assets. Lai, Xing, and Chen (2011) develop an approach for optimizing portfolios based on their mean-variance when the means and covariances are not known. Yu and Yuan (2011) investigate how investor sentiment influences the mean variance tradeoff while Palczewski and Palczewski (2014) develop an approach of estimating the sensitivity of mean-variance portfolios.

A few studies have used the portfolio theory framework to investigate asset trading. Anderson (2013) investigates the relationship between trading and diversification while Gârleanu and Pedersen (2013) build upon Markowitz's approach to develop an optimal portfolio policy when trading is dynamic. Zhao and Palomar (2018) also extend Markowitz's approach to study the optimal portfolio investment with options trading.

International portfolio investment is another area that has been greatly influenced by modern portfolio theory. In particular, the process of diversification proposed by Markowitz (1952) has been of interest. Levy and Sarnat (1970), Rugman (1976), Brewer (1981), Levy and Lerman (1988), Campa and Fernandes (2006), Dreissen and Laeven (2007), and Coeurdacier and Guibaud (2011) all illustrate benefits and gains investors can earn by diversifying in the international market. Similar studies have extended these applications to more specific geographical markets as well. Lessard (1973) performs an analysis of international diversification in Latin American countries while Eun and Resnick (1994) focus on the United States and Japan. Gilmore and McManus (2002) and Syriopoulos (2004) conducts similar analyses for the Central European markets.

The portfolio theory framework has been utilized in several other industry specific studies within finance as well. Friedman (1971) applies the Markowitz approach to the process of selecting real estate portfolios. Arouri and Nguyen (2010) study the relationship between oil prices and the stock market and show that investors can benefit from including an oil asset in an already well-diversified portfolio of stocks. Brauneis and Mestel (2019) use the mean-variance framework to analyze the risk and returns of cryptocurrency portfolios. Platanakis and Urquhart (2020) apply the framework to address whether investors will benefit from including bitcoin in their portfolios.

The portfolio theory framework has clearly had a tremendous impact on a variety of areas within the finance literature. However, more broadly, this framework has extended far beyond the financial literature and into the research of several other disciplines. Firm investment is one such area in which a wide variety of portfolio theory applications have been used. Cardozo and Smith (1983) show how portfolio theory from finance can be applied to a firm's product portfolio decisions. Lubatkin and Chatterjee (1994) use portfolio theory to investigate the relationship between corporate diversification and risk. Yorke and Droussi-otis (1994) provide examples of how the framework can be useful in marketing by treating

customers as assets that can be defined within a customer portfolio. Hickman, Teets, and Kohls (1999) bring the dimension of a firm's social responsibility into the portfolio theory framework and find that socially responsible firms can help to reduce risk in a portfolio. Choi, Li, and Han (2008) use a mean-variance approach to analyze the newsvendor problem in operations and inventory management. The approach allows the authors to account for risk preferences in the process in addition to the expected cost or profit. Wu et al. (2009) extend this mean-variance approach to the newsvendor problem to account for stockout costs in the model. Teller and Kock (2013) apply the portfolio theory framework to a project management setting by investigating the link between managing the risk of a portfolio of projects and the success of the project portfolio.

Other areas in which the portfolio theory framework has been applied include imports, energy, and biology. Wu et al. (2007) use portfolio theory to investigate the risk of crude oil imports in China by quantifying the diversification index of these imports. Roques, Newbery and Nuttall (2008) find value in portfolio theory by applying the mean variance framework to the electricity market to find the optimal portfolio choice for electricity generation. Nagengast, Braun, and Wolpert (2011) use the mean-variance framework to analyze decision making in the sensorimotor system. The movements observed in completing sensorimotor system tasks are consistent with a risk-return trade-off. Zhang, Zhao, and Xie (2018) use portfolio theory as an approach to developing the optimal mix of power generation in China.

As evidenced by the previous literature, the impact of portfolio theory has been far-reaching both within its original discipline of finance and beyond into a wide range of other disciplines. However, one potential application of portfolio theory that has yet to be demonstrated in the academic literature pertains to the construction and performance of a team. In recent years, the areas of sport economics and sport finance have grown considerably in the academic literature as researchers have found value in using the sports industry to empirically test a wide range of theory pertaining to economics and finance. One of the biggest

benefits to using sports data in these disciplines is that employee (player) productivity can be observed and objectively quantified, something that is rare and difficult to do in most other industries.

A professional sports team is a setting for which the portfolio theory framework can be extended quite naturally. A sports team consists of a collection of players, each of which contribute to the success of the team through their productivity. This can be viewed as analogous to a financial portfolio, which is defined by a collection of financial assets, each of which contribute to the overall performance of the portfolio. Thus, we can think of a sports team as a portfolio in which the players represent the assets. Furthermore, given the rich data sources available on player productivity, we can define player productivity within a mean-variance framework such that each player contributes to the mean, variance, and ultimately performance of their respective teams.

This paper addresses two primary gaps in the current literature. The first is the extension of the portfolio theory framework to a sports team such that teams can be defined as a portfolio of assets characterized by its mean performance level and its performance variance. Using this portfolio theory application, we then perform an analysis of efficiency in the player trade market in Major League Baseball. This is the second major gap in the research this paper seeks to address, as there has yet to be a great deal of research on the trade market in the MLB. Ultimately, the applications of portfolio theory used in this paper can be extended to any team setting in which member productivity can be reliably measured.

3.3   Run, Run (Expectancy), Run (Value)

A batting event in baseball is defined as the outcome of a plate appearance which either advances one or more players around the bases, or results in an out. Every time a batter comes to the plate, there are five principal ways he can contribute to scoring runs: a walk, a

single (both worth one base) a double (two bases), a triple (three bases) or a home run (all four bases).[1] These five types of event account for around 30% of all events in a season of Major League Baseball (there are around 200,000 events each season). Around two thirds of events result in an out.[2]

Since the earliest days of professional baseball, there have been attempts to quantify the contribution of a batter to the runs scored by the team, for the purpose of ranking performance. The concept of batting average (hits divided by at bats) goes back to Henry Chadwick, the Father of Baseball, who proposed the measure in the year 1867.[3] Slugging percentage developed as a way of crediting batters for hits which gained more bases. From the 1970s onwards analysts such as Bill James developed extended notion such as on-base-percentage (OBP) and On-base-plus-slugging (OPS) as richer measures of batting contribution. The greatest technical advance in the measurement of batting performance developed out of the work of Lindsey (1963).

3.3.1 Run Expectancy and Run Value

The intuition behind the run expectancy matrix is that batting statistics are context dependent – the value of an action depends on the state of the half inning in which it occurs. Whenever a batter steps up to the plate, the half inning is in one of 24 potential states, which are defined by the states of the bases and the number of outs – see Figure 3.1.

---

[1]As well as advancing the batter around the bases, these events can also advance previous batters who are currently on base. In the extreme, if there are runners on each base (bases loaded) then a home runs results in four runs scored in total.

[2]A relatively small fraction of events can both advance players on base and cause an out- these are mostly sacrifice plays (about 2% of all events). Hit by pitch (which advances the batter to first) and fielding errors, each accounting for about 1% of events.

[3]https://ourgame.mlblogs.com/chadwicks-choice-the-origin-of-the-batting-average-e8e9e9402d53

Figure 3.1: The Run Expectancy Matrix

	Outs		
	0	**1**	**2**
000			
100			
010			
Base **110**			
States **001**			
101			
011			
111			

Thus base state 000 means no runners on base, 100 means a runner on first and no runners on second or third, 010 means no runner on first, a runner on second and no runners on third, and so on. The run expectancy matrix is then populated by finding the average number of runs scored from any given state to the end of the half inning, for all events in an entire season. This can be done using event data available at Retrosheet.org or in recent years from files made available by MLBAM.[4]

Table 3.1: The Run Expectancy Matrix for the 2011 MLB Season

		Outs		
		0	**1**	**2**
	000	0.471	0.255	0.097
	100	0.835	0.496	0.218
Base	**010**	1.058	0.650	0.309
States	**110**	1.414	0.874	0.422
	001	1.454	0.937	0.317
	101	1.753	1.150	0.488
	011	1.930	1.339	0.541
	111	2.172	1.475	0.761

---

[4]https://www.rdocumentation.org/packages/mlbgameday/versions/0.2.0

Table 3.1 is the run expectancy matrix for the 2011 MLB season, based on 191,864 events. Thus, for example, the average number of runs scored, from the state of no outs and no runners on base until the end of the half inning, was 0.471. The number of runs scored on average with no runners on base and two outs until the end of the half inning was 0.097. Not surprisingly, run expectancy tends to increase when there are more runners on base, when the runners are further advanced, and when there are fewer outs. There are some variations in the run expectancy matrices from year to year, but generally these variations are quite small. Over the timeframe 1994-2016, the standard deviation in run expectancy for the season was below 0.05 runs for 17 out of the 24 base states and ranged from 0.008 runs to 0.102 runs overall. As expected, the rarer, less frequent base states (e.g. bases loaded and no outs) resulted in the most variation in run expectancy from year to year.

From the run expectancy matrix we can calculate the average run expectancy of a given event (a single, a walk, a home run, an out, etc.) in a season. For any event, there is a change in run expectancy which is defined as the difference between the run expectancy immediately before the event (the start state) and run expectancy when the event ends (the end state).[5] For example, for the 2011 season, for any given event, the change in run expectancy can be derived simply by taking the difference in cell values, identifying the start state and the end state in Table 3.1. The run value of an event can then be defined as the runs scored during the event, plus the difference in run expectancy at the beginning of an event and at the end of the event:

---

[5]Technically, to define all possible run values we need to define the run value of end states where there are three outs (since these end states occur). For each of the eight possible states (relating to the configuration of the base runners) the run expectancy is zero, since the half inning is over.

44

$$Run\ Value = Runs\ Scored\ During\ Event\ +$$

(3.1)
$$Run\ Expectancy\ at\ the\ End\ of\ the\ Event\ -$$

$$Run\ Expectancy\ at\ the\ Beginning\ of\ the\ Event.$$

Using the empirical frequencies we can construct a table of run values for a given event, depending on the starting states. This is illustrated for the home run event in the 2011 season in Table 3.2.

Table 3.2: Run Value of Home Runs in 2011

	0	1	2
000	1.000	1.000	1.000
100	1.636	1.759	1.879
010	1.413	1.604	1.788
110	2.057	2.381	2.675
001	1.017	1.317	1.780
101	1.718	2.105	2.609
011	1.541	1.916	2.556
111	2.299	2.780	3.336

The run value of a home run, with no one on base at the beginning of the event, is always 1, which is just the run scored, given that there is no one on base at the end of the event (i.e. the bases states are unchanged). At the other extreme, the run value of bases loaded and two outs (111 2) is 3.336 runs, which is the value of the four runs scored plus the change in value of the base states.[6]

---

[6]At first glance, this might seem counterintuitive- four runs were scored, but the run value is less than four. The intuitive explanation is that while four runs were scored, the event itself should not be given the full credit for these runs, since hitting the home run was not responsible for getting the three runners on base at the beginning of the event – that achievement should be credited to earlier events. Note that at the beginning of the event, when bases were loaded, the run expectancy was very high, but at the end of the event there are no runners on base, so that run expectancy has fallen considerably.

We can also calculate the average value of an event across all starting states. Thus, for example, the run value of a home run, in a given season, is the weighted sum across starting states of the run values of a home run, where the weights are the frequencies of each starting state for that event. Based on the run values in Table 3.2 and the frequency of home runs from each starting state in 2011, the average run value of a home run was 1.39 in that season.

Table 3.3 shows the weighted run values of each of the six main batting events for each of the seasons in our data. These values are often referred to as "linear weights" in the literature.[7] It is clear from Table 3.3 that these values are very stable over time. This is not the same as saying that the game never changes. Over this period there have clearly been many innovations in baseball, not least those associated with *Moneyball* and the increasing use of data analytics. The first half of the data also roughly coincides with the "steroid era", which increased batting power and thus the production of home runs. While the amount of production of particular events may have changed over time, what Table 3.3 indicates is the frequency with which particular events are converted into runs has remained very stable over this period.

---

[7]Albert and Bennett (2001) attribute this coinage to Peter Palmer (see p. 200).

Table 3.3: Run Value Weights by Season for Batting Events

Season	Walk	Single	Double	Triple	Home Run	Out
1994	0.299	0.467	0.781	1.064	1.400	-0.289
1995	0.319	0.482	0.767	1.074	1.405	-0.291
1996	0.331	0.485	0.784	1.105	1.403	-0.302
1997	0.307	0.465	0.761	1.083	1.393	-0.284
1998	0.312	0.469	0.780	1.014	1.400	-0.285
1999	0.311	0.477	0.789	1.059	1.408	-0.302
2000	0.332	0.482	0.765	1.085	1.406	-0.307
2001	0.298	0.460	0.778	1.084	1.380	-0.283
2002	0.303	0.466	0.755	1.052	1.398	-0.279
2003	0.307	0.466	0.775	1.080	1.391	-0.284
2004	0.307	0.462	0.786	1.041	1.396	-0.287
2005	0.295	0.458	0.768	1.056	1.412	-0.277
2006	0.317	0.467	0.766	1.070	1.389	-0.290
2007	0.310	0.468	0.798	1.044	1.406	-0.289
2008	0.312	0.460	0.772	1.081	1.405	-0.281
2009	0.304	0.459	0.762	1.004	1.392	-0.278
2010	0.299	0.451	0.763	1.076	1.404	-0.266
2011	0.289	0.442	0.736	1.064	1.392	-0.255
2012	0.284	0.441	0.747	1.039	1.396	-0.257
2013	0.285	0.439	0.740	1.035	1.371	-0.250
2014	0.283	0.437	0.739	1.054	1.400	-0.245
2015	0.303	0.442	0.743	1.031	1.386	-0.257
2016	0.288	0.439	0.739	1.014	1.382	-0.264
mean	0.304	0.460	0.765	1.057	1.396	-0.278
standard deviation	0.014	0.015	0.018	0.026	0.010	0.017

3.4 Team Portfolios

Just as financial portfolios are comprised of a collection of assets, each of which contributes to the total returns of the portfolio, a sports team can be thought of as a portfolio of players, each of which contributes to the total team production. Furthermore, team portfolios can be characterized by the mean and variance of team production based on the individual performance means and variances of its players. In baseball, each player's production can be characterized by his expected (mean) run value contributions during a game and the variance of this measure.

To calculate each player's expected run value and variance, we compute the sum of each player's run value contributions at the game level for every game in which the player received a plate appearance during the season. Thus, these game level run value contributions will represent the set of "daily returns" for each player in the season for which we will use to compute the player's expected game level run value and variance. We compute these measures using Retrosheet's Event Logs for all batting events from 1994-2016.[8] The distributions of mean and variance for all players with at least 30 games with a plate appearance during the season are included below.

---

[8]https://www.retrosheet.org/game.htm

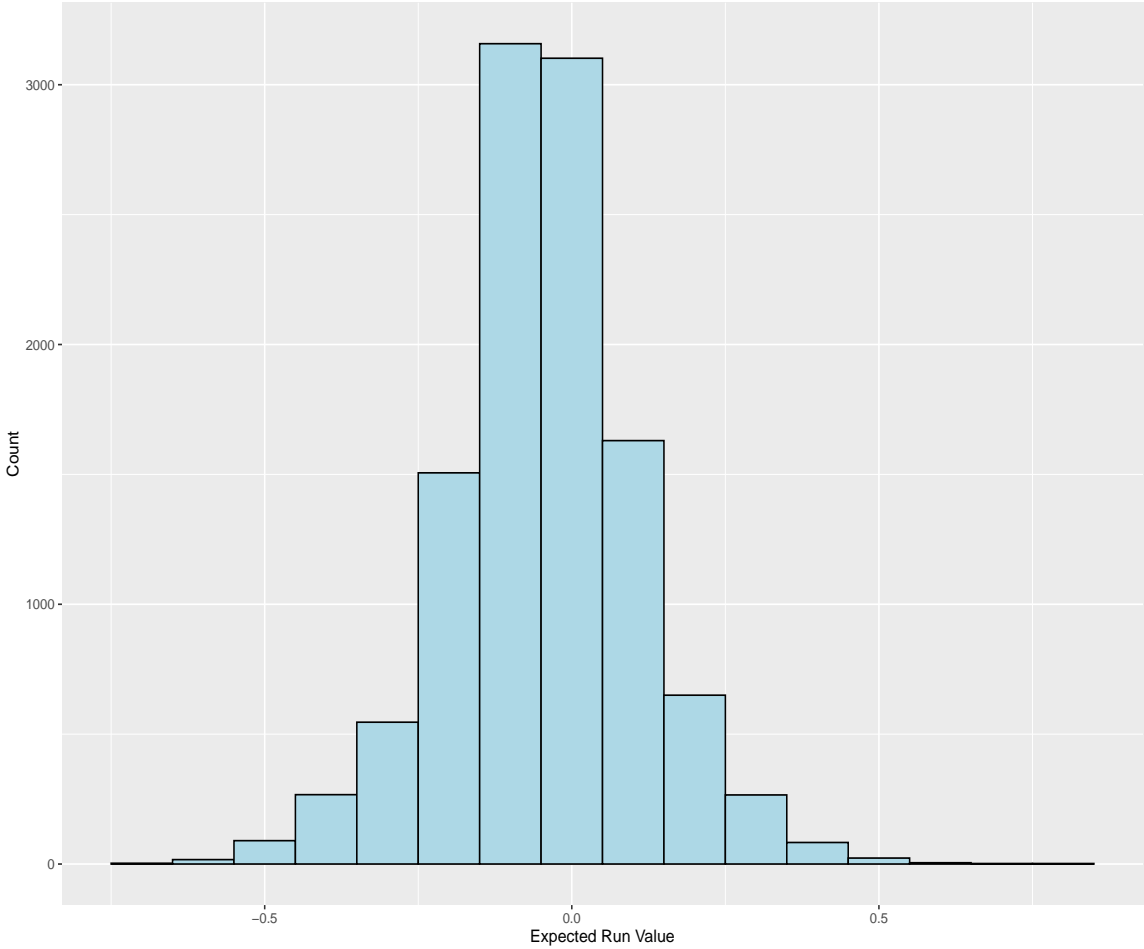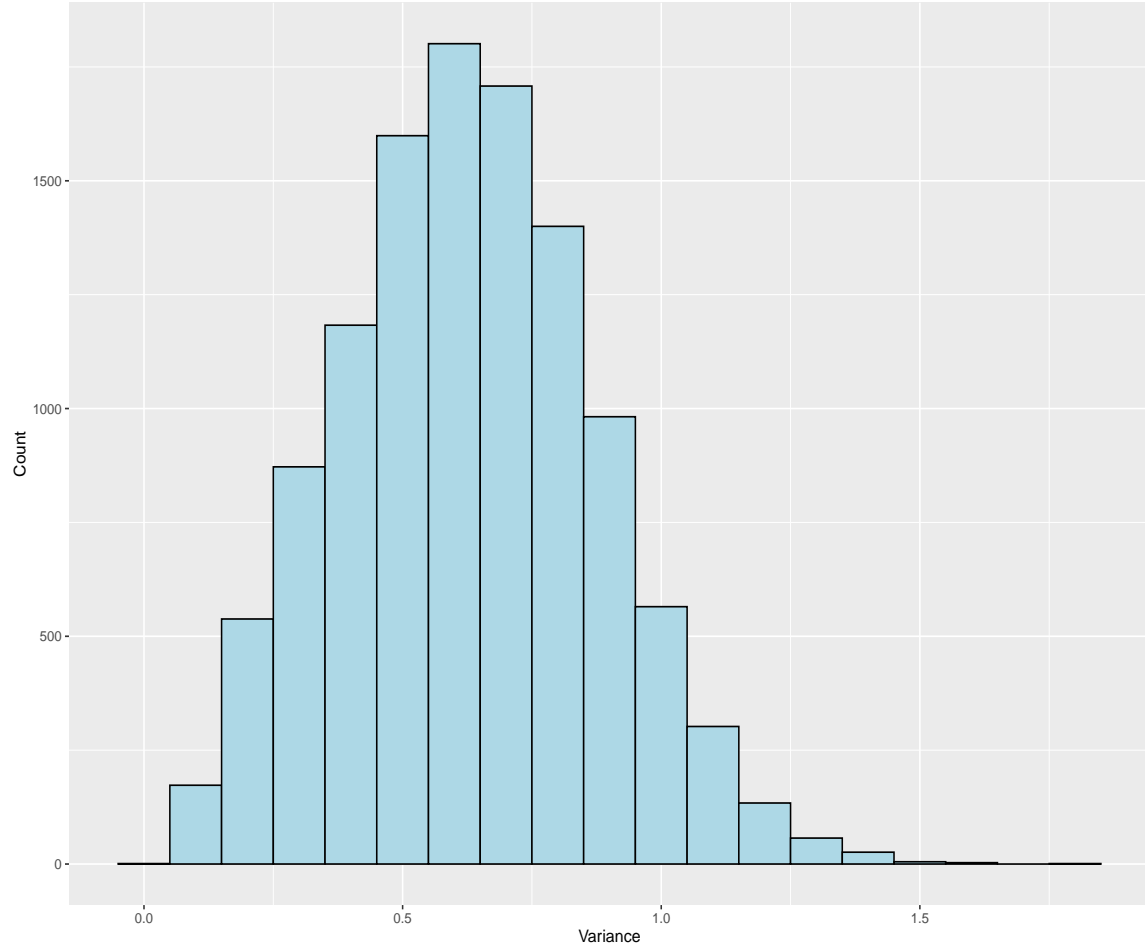Figure 3.2: Histogram of Player Expected Run Values: 1994-2016

Figure 3.3: Histogram of Player Variances: 1994-2016

Using each player's mean and variance, we can calculate team level measures of mean and variance using the standard approach for computing the mean and variance of a financial portfolio. In financial portfolio theory, the mean and variance of a portfolio's returns are calculated as follows:

$$(3.2) \qquad Expected\ Portfolio\ Return: E(R_p) = \sum_{i}^{n} w_i E(R_i)$$

$$(3.3) \qquad Portfolio\,Variance: Var(R_p) = \sum_i^n \sum_j^n w_i w_j Cov(R_i, R_j),$$

where R_i generally represents the set of daily returns for each asset i in the portfolio, n represents the number of assets in the portfolio and w_i represents the share of the total portfolio value belonging to asset i.

Using this framework, we can treat each team as a portfolio of players with individual means and variances based on their game level run value contributions. Therefore, the game level expected (run) value and variance of a given team during the season can be calculated as follows:

$$(3.4) \qquad Expected\,Team\,Performance: E(TPERF_k) = \sum_i^{n_k} w_i E(PERF_i)$$

(3.5)

$$Variance\,of\,Team\,Performance: Var(TPERF_k) = \sum_i^{n_k} \sum_j^{n_k} w_i w_j Cov(PERF_i, PERF_j)$$

where $E(PERF_i)$ represents the mean run value per game of each player i on team k, n_k represents the number of players on team k that received at least one plate appearance in the season and w_i represents player i's share of total team games played for team k over the course of a season. For the team variance calculation, when $i = j$, the right hand side of the formula reduces to $\sum_i^{n_k} w_i^2 Var(PERF_i)$, where $Var(PERF_i)$ represents the variance of player i's game level run value contributions.

However, when $i \neq j$, we need to calculate measures of covariance for each pairwise

combination of batters on a team for a given season. To do this, we match up the game level run value contributions for every pairwise combination of batters on a team for the season. As a result, only games in which both teammates receive at least one plate appearance are included in the data for our game level run value covariance calculations for pairwise combinations of teammates. For any two teammates i and j on team k, the covariance can be calculated as follows:

$$(3.6) \quad Cov(PERF_i, PERF_j) = \frac{1}{n_g - 1} \sum_{g=1}^{n_g} (PERF_{ig} - E(PERF_i))(PERF_{jg} - E(PERF_j))$$

In the above equation $PERF_{ig}$ and $PERF_{jg}$ represent the actual game level run value contributions for players i and j in game g. $E(PERF_i)$ and $E(PERF_j)$ represent the mean game level run value contributions for players i and j for the season as calculated in Equation (3.4) and n_g represents the total number of games during the season for which players i and j both received at least one plate appearance. Using this approach we calculate the covariance for every pairwise combination of batters on a team in our data and then we use these covariance calculations to calculate team performance variance as shown in Equation (3.5).

Using our measures of the game level expected (run) value and variance of a given team's player during the season, we run an OLS regression of team runs for the season on these measures of mean and variance. To account for differences between seasons, we use the z-scores of the team level means and variances standardized by season. The formal model and regression results are as follows:

$$(3.7) \qquad ZRuns = \beta_0 + \beta_1 ZMean + \beta_2 ZVar + \sum_i \alpha_i Team_i + \epsilon$$

Table 3.4: Contribution of Team Mean and Variance to Runs

	Coefficients
ZMean	0.867***
	(0.0175)
ZVar	0.107***
	(0.0187)
Constant	0.203**
	(0.0701)
R^2	0.905
Adj. R^2	0.900
Num. obs.	682

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

From our regression results we observe that the game level expected run value and variance for a team explains the majority of variation in total team run production with an R-squared of .905. We also observe that both mean and variance are positive and highly significant, though a team's expected run value is more closely tied to team runs than a team's variance. Given this result, a reasonable follow up question may to be ask why variance is a positive and significant contributor to team run production when holding the team's mean level of production constant. One might expect the impact of variance to be insignificant (or even negative) when comparing teams with similar expected run values. To help address this question, we propose the following thought experiment.

Suppose we have two batting lineups that consist of nine identical batters. The first lineup has batters who only reach base via singles, each with the same probability of success. The second lineup has batters who only reach base via home runs, each with the same probability of success. Based on our definition of variance, the singles only lineup consists of lower variance players while the home runs only lineup consists of higher variance players.

In the case of the "singles only" lineup, we will assume that each success (single) advances the batter and any baserunners by one base. Therefore, the number of runs scored in an inning for the singles only lineup will be as follows, where X represents the number of batters

that come to the plate in an inning:

$$Runs\ Scored\ (Singles\ Only) = \begin{cases} 0,\ 3 \leq X < 7 \\ X - 6,\ X \geq 7 \end{cases}$$

For a team to score at least one run in this scenario, X must be at least 7 and the number of successes must be at least 4 (since 4 singles result in one run scored). Any value below 7 will not result in a run scored since the plate appearances of 3 batters in the inning result in an out and the number of successes will be less than 4.

In the case of the "HR only" lineup, every success will result in one run scored by definition of a home run. Thus, the number of runs scored in an inning for the HR only lineup will be:

$$Runs\ Scored\ (HR\ Only) = \begin{cases} 0,\ X = 3 \\ X - 3,\ X \geq 4 \end{cases}$$

These simple production processes provide us with some intuition into why higher variance players tend to add value to teams compared to lower variance players when controlling for the mean. In the singles only (lower variance) case, a greater degree of sequential production (worker coordination) is required for the team to score runs. Given that each success in the singles only lineup advances each baserunner by one base, the successes of teammates need to be sequentially clustered in order for the team to score any runs. Therefore, the order in which the team achieves its successes is a major factor in team production for the singles only lineup. In the HR only case, a success results in a run by definition and no co-ordination between teammates is needed. Thus, contrary to what we observe in the singles only lineup, the sequence in which teammates achieve their successes is not a relevant factor in team production for the HR only case. Therefore, the advantage teams gain for having higher variance while controlling for mean arises from the fact the run-scoring process for

the higher variance lineup requires less dependence between workers than the run-scoring process for lower variance lineups.

3.5   Trade Efficiency

We now extend the framework developed in the previous section to develop an approach to analyze trade efficiency between teams. The results from our previous section show that players add runs to their team through mean and variance and each team consists of a portfolio of players whose combined means and variances deliver runs across a season. When considering the talent distribution for players, we also anticipate a certain trade-off between player variance and mean level of performance, as we expect variance to be more valuable for lower mean levels due to the sequential nature of team production in baseball.

We can define an isoruns curve as the level set of combinations of mean and variance that will deliver the same number of runs across a season. Thus the isoruns curves is defined by:

$$(3.8) \qquad\qquad dR = \frac{\partial R}{\partial \mu}d\mu + \frac{\partial R}{\partial \sigma^2}d\sigma^2 = 0.$$

We parameterize the model as:

$$(3.9) \qquad\qquad R = \mu^\alpha(\sigma^2)^\beta,$$

which we can then estimate as:

$$(3.10) \qquad\qquad \ln R = \alpha \ln \mu + \beta \ln \sigma^2,$$

to recover α and β.

The slope of any isoruns curve is defined as:

$$(3.11) \qquad \frac{d\mu}{d\sigma^2}|_{R=\bar{R}} = -\frac{\frac{\partial R}{\partial \mu}}{\frac{\partial R}{\partial \sigma^2}} = -\frac{\frac{\alpha R}{\mu}}{\frac{\beta R}{\sigma^2}}.$$

From this derivation we make the following proposition:

Proposition 1. An efficient trade requires that the gain in value to the buying team from hiring player i should be greater than the loss in value to the selling team from selling player i, so G_{Bi}–L_{Si}, where

$$(3.12) \qquad G_{Bi} = dR_{Bi} \approx \frac{\alpha R_B}{\mu_B}\mu_i + \frac{\beta R_B}{\sigma_B^2}\sigma_i^2$$

$$(3.13) \qquad L_{Si} = dR_{Si} \approx \frac{\alpha R_S}{\mu_S}\mu_i + \frac{\beta R_S}{\sigma_S^2}\sigma_i^2.$$

To assess the efficiency of trades over our timeframe we measure these quantities for each trade and determine what proportion of trades satisfy this condition.

In the above equations, R_B, μ_B, and σ_B^2 represent the buying team's runs per game, mean, and variance prior to the trade respectively. The variables R_S, μ_S, and σ_S^2 represent the same measures for the selling team. Finally, μ_i, and σ_i^2 represent the mean and variance of the traded player i, prior to the trade. To determine whether a trade is efficient, we calculate the quantity G_{Bi}–L_{Si} and classify it as an efficient trade if this value is positive. We then use these results to analyze trade efficiency across the MLB.

We begin by estimating our values of α and β using the following regression model:

$$(3.14) \qquad lnRuns = \gamma + \alpha lnMeanTrans + \beta lnVar + \epsilon.$$

In this regression, $lnRuns$ is the natural log of team runs per game for the season. The variable $lnMeanTrans$ is the natural log of each team's mean transformed by adding the absolute value of the minimum team mean to all observations prior to taking the natural log so that all values are defined once taking the natural log. Finally, $lnVar$ represents the natural log of each team's variance of the season.

After estimating the appropriate exponents for our production function, we construct isorun curves over the range of team runs per game using the values of α and β estimated in the regression. Since our timeframe of 1994-2016 covers a wide range of years, we construct separate isorun curves for 5 different periods within this timeframe since the impact of mean and variance on team run production is unlikely to stay constant.

To determine the breakpoints for our analysis, our regression was run for all possible combinations of time periods ranging for 4-7 years. We then selected breakpoints based on the set of time periods resulting in the highest average R-squared value across eras. The periods used and corresponding R-squared values are summarized in the table below.

Table 3.5: R^2 by Era

| Time Period	R^2
1994-1997	0.862
1998-2001	0.819
2002-2005	0.878
2006-2009	0.850
2010-2016	0.810

Our isorun curves are constructed for each respective era with the transformed team mean on the horizontal axis and team variance on the vertical axis. Included in these figures are also the scatterplot points of team transformed mean and team variance for all teams over the different time periods. This allows us to visualize where along the isorun curves the team portfolios tend to lie. Each curve represents a level of team production in the form of runs per game, and each curve represents the mean-variance tradeoff required for teams to maintain the same level of production. The regression results used to estimate the values of α and β for each era and the resulting isorun curve plots are included below as well.

Table 3.6: Isorun Curve Coefficients

	1994-1997	1998-2001	2002-2005	2006-2009	2010-2016
logMeanTrans	1.999***	1.952***	2.285***	2.045***	2.255***
	(0.114)	(0.129)	(0.107)	(0.112)	(0.115)
logVar	0.0603**	0.0832***	0.0661***	0.0667***	0.0689***
	(0.0186)	(0.0188)	(0.0140)	(0.0143)	(0.0131)
Constant	1.607***	1.687***	1.506***	1.605***	1.415***
	(0.0611)	(0.0625)	(0.0488)	(0.0474)	(0.0497)
R^2	0.862	0.819	0.878	0.850	0.810
Adj. R^2	0.860	0.816	0.876	0.847	0.808
Num. obs.	112	120	120	120	210

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

Figure 3.4: Isorun Curves by Era (With Transformed Means)
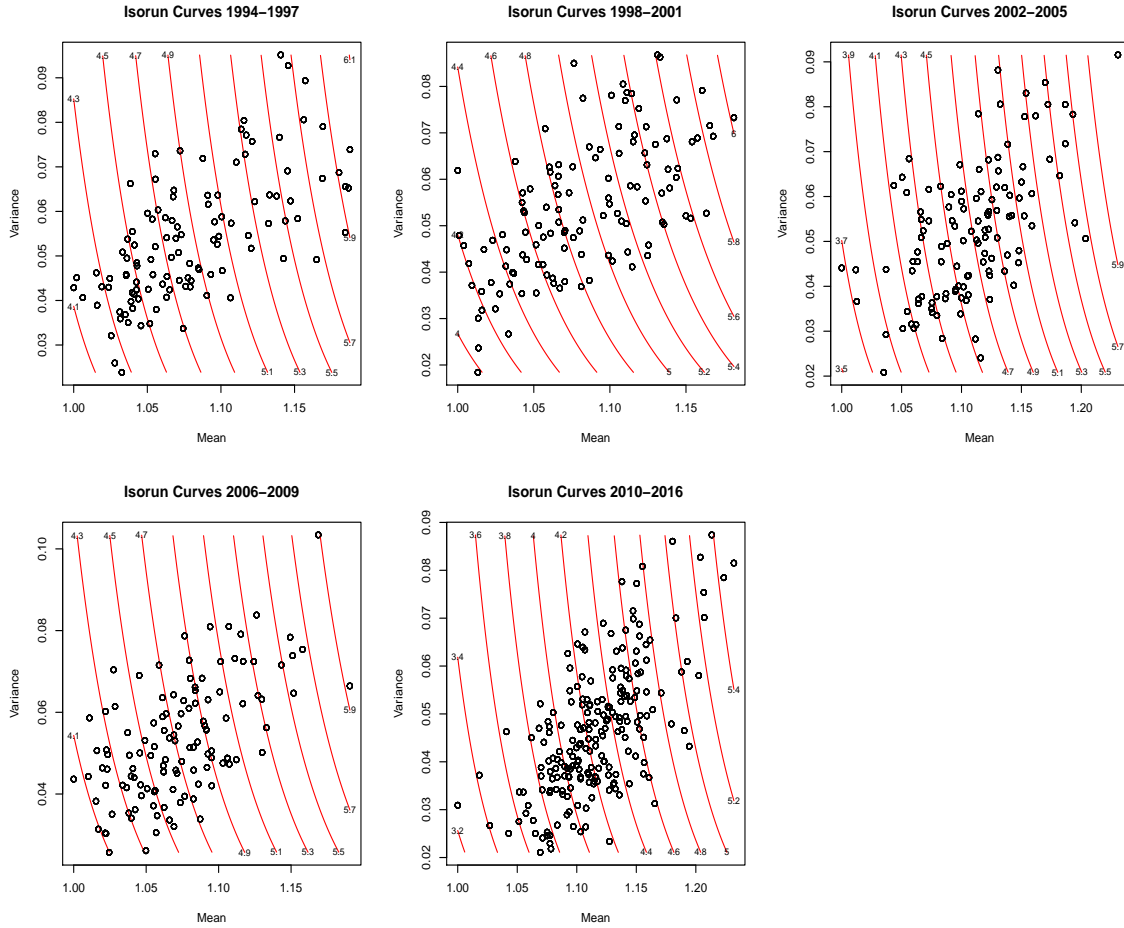
From the plots and the regression results we do observe some differences in the contributions of team mean and variance to runs per game by era. However, in each era mean and variance are positive and significant with mean being the bigger contributor to team runs per game. These results match the findings from our results in the previous section.

Having established teams as portfolios of players in which returns (team production) can be defined as a function of team means and variances, we would now like to use this framework to investigate trade efficiency between teams over time. Using transaction data from Retrosheet, we focus on trades for hitters that occurred in the middle of the season

59

(May, June, July or August).[9] We exclude offseason trades and trades that take place at the very beginning of the season so that we have a reliable measure of team offensive quality at the time the trade occurs.

To be included in our analysis the traded players must have had at least 100 PA for the selling team prior to the trade and at least 100 PA for the buying team after the trade occurs during the season. These constraints allow each player's performance to stabilize prior to the trade while ensuring that the buying team intends on giving the traded player regular playing time after the trade occurs. Dating back to 1994, we identify 364 transactions in which these conditions are met.

For each transaction, we calculate the net gain using the approach in Proposition 1 to determine whether the transaction was efficient. We then calculate and plot the efficiency rate by season to observe if any changes to trade efficiency have occurred over time. We exclude 1994 from our plot since there were only 3 trades that met our criteria this season due in large part to the strike that ended the season in August. The resulting plot for 1995-2016 is as follows.

---

[9]https://www.retrosheet.org/transactions/

Figure 3.5: Trade Efficiency Rate by Season: 1995-2016
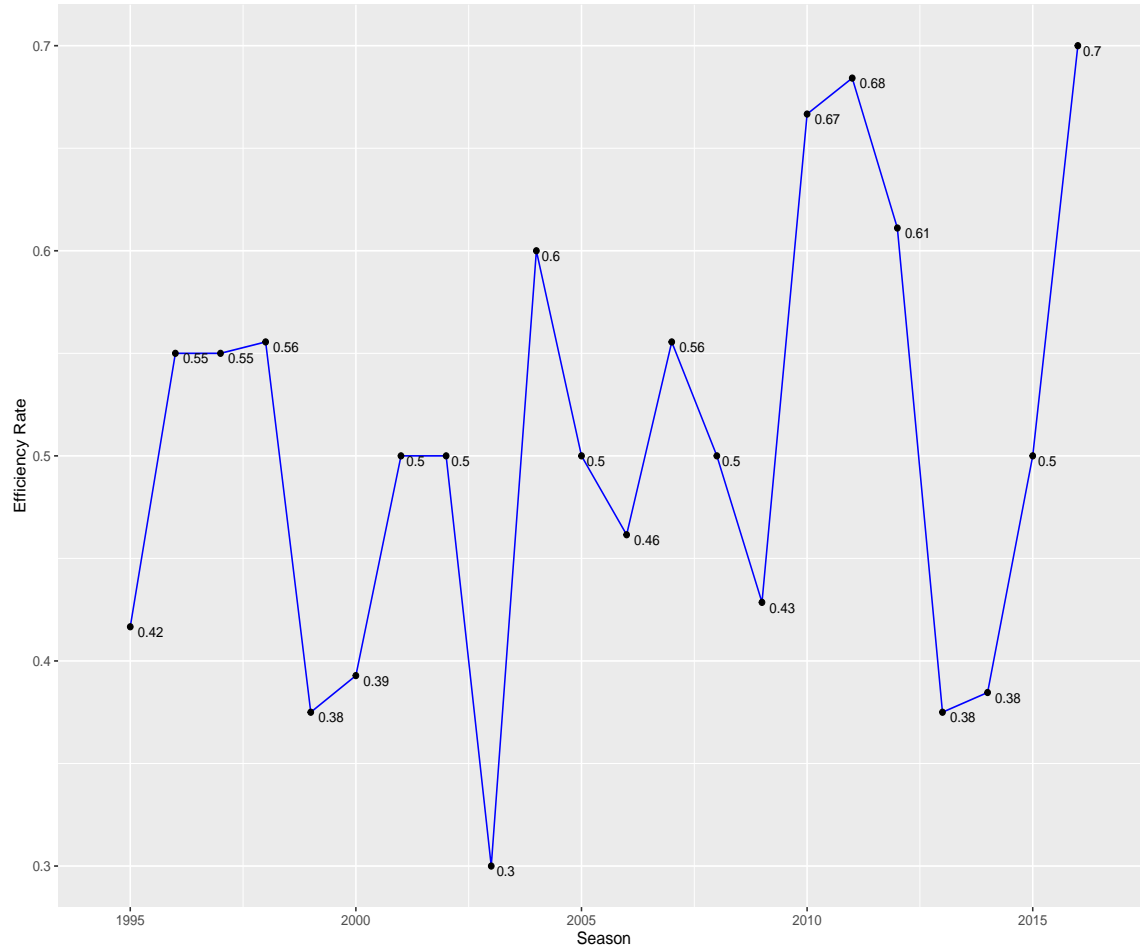
When aggregating trades over the entire timeframe, we calculate an efficiency rate of 50.5%. We also observe a general increasing trend in efficiency rate over time. Four out of the five highest efficiency rates at the season level take place from 2010 onwards with the maximum efficiency rate reaching as high as 70% in the most recent season. When aggregated, the most recent era of 2010-2016 has an efficiency rate of 57.9%, 5 percentage points higher than the next highest efficiency rate by era where eras are defined as in Table 3.6. Intuitively these findings are consistent with the increase in analytics and data driven decisions and less emphasis on traditional scouting used by MLB front offices over this

timeframe.

These can be viewed as conservative estimates of trade efficiency since trades are often much more complex than the subset of trades that meet our criteria. Thus, based on our definition of efficiency, an inefficient trade does not necessarily imply a poor decision was made. The following are potential examples of such cases: (1) The buying team may have a pressing organizational need at a specific position (i.e. due to player injury) and may be willing to pay a premium for the acquired player. (2) The selling team may be trading from a position of organizational depth (the team may already have a viable alternative to the player being traded away on the roster or in the minor leagues). (3) The selling team may be more interested in acquiring talent with a greater future value than the player being sold (these transactions usually involve young, unproven talent often from the minor leagues). In addition, recall that we are excluding all pitching transactions from this analysis as well.

Thus, if one were able to account for some of these complexities involved in the decision-making process regarding player transactions, it is possible that this efficiency rate would increase. However, despite these complexities, we are still able to use our framework to observe the overall trend of trade efficiency over time. We observe from our results that improvements in trade efficiency over time are consistent with teams making more informed, data-driven decisions over the course of our timeframe.

Out of the 364 trades identified in our analysis, 184 of these were classified as efficient by our definition of efficiency. As a robustness check, we use a Monte Carlo simulation to determine the probability of each trade remaining efficient when our transformed mean and variance parameters are subject to random shocks at both the team and player level. Thus, we would like to re-estimate our values of G_{Bi} and L_{Si} as follows:

(3.15) $$G_{Bi} = dR_{Bi} \approx \frac{\alpha R_B}{\mu_B + \epsilon_{\mu_B}}(\mu_i + \epsilon_{\mu_i}) + \frac{\beta R_B}{\sigma_B^2 + \epsilon_{\sigma_B^2}}(\sigma_i^2 + \epsilon_{\sigma_i^2})$$

(3.16) $$L_{Si} = dR_{Si} \approx \frac{\alpha R_S}{\mu_S + \epsilon_{\mu_S}}(\mu_i + \epsilon_{\mu_i}) + \frac{\beta R_S}{\sigma_S^2 + \epsilon_{\sigma_S^2}}(\sigma_i^2 + \epsilon_{\sigma_i^2})$$

For μ_B, μ_S, σ_B^2, and σ_S^2 the error terms are obtained from the regressions used to estimate our parameters in the team run production function for each respective era. Each regression is reconstructed to solve for the errors in terms of $lnMeanTrans$ and $lnVar$ respectively. Then, ϵ_{μ_B}, ϵ_{μ_S}, $\epsilon_{\sigma_B^2}$, and $\epsilon_{\sigma_S^2}$ are randomly drawn from these respective vectors of error terms and substituted into our equations for G_{Bi} and L_{Si}.

For μ_i and σ_i^2 the error terms are obtained from the following regression at the player level and run separately over the same eras as defined in the team run production analysis:

(3.17) $$lnPlyrRuns = \gamma + \alpha lnMeanTransPlyr + \beta lnVaPlyr + \epsilon.$$

In this regression, $lnMeanTransPlyr$ represents the natural log of each player's mean transformed by adding the absolute value of the minimum player mean to all observations so that so that all values are defined once taking the natural log, while $lnVarPlyr$ represents the natural log of each player's variance. To match the approach we have used to model team productivity under the mean-variance framework, our dependent variable in this regression will be the natural log of player runs per game, simply equal to the number of runs scored by the player during the season divided by the number of games played. As in our team level regressions, each regression is reconstructed to solve for the errors in terms of

$lnMeanTransPlyr$ and $lnVarPlyr$ respectively. Finally, ϵ_{μ_i} and $\epsilon_{\sigma_i^2}$ are randomly drawn from these vectors of error terms and substituted into our equations for G_{Bi} and L_{Si}.

For each trade that we identified as efficient, we run 10,000 simulations recalculating the values of G_{Bi} and L_{Si} while subjecting μ_B, μ_S, μ_i, σ_B^2, σ_S^2, and σ_i^2 to random shocks. For each simulation, these shocks are randomly drawn from the appropriate error vectors and used to calculate the net gain of the trade $G_{Bi} - L_{Si}$. Finally, for each trade we calculate the proportion of simulations in which the net gain remains positive. This serves as an estimate for the probability the trade is efficient given random shocks to our parameters.

For each of our 184 efficient trades, we find that the average probability of efficiency when subject to random parameter shocks is 63.84%. As an illustration, we plot each trade's probability of efficiency as a function of the trade's net gain (without being subject to random shocks).

Figure 3.6: Net Trade Value Gain vs. P(Efficient)
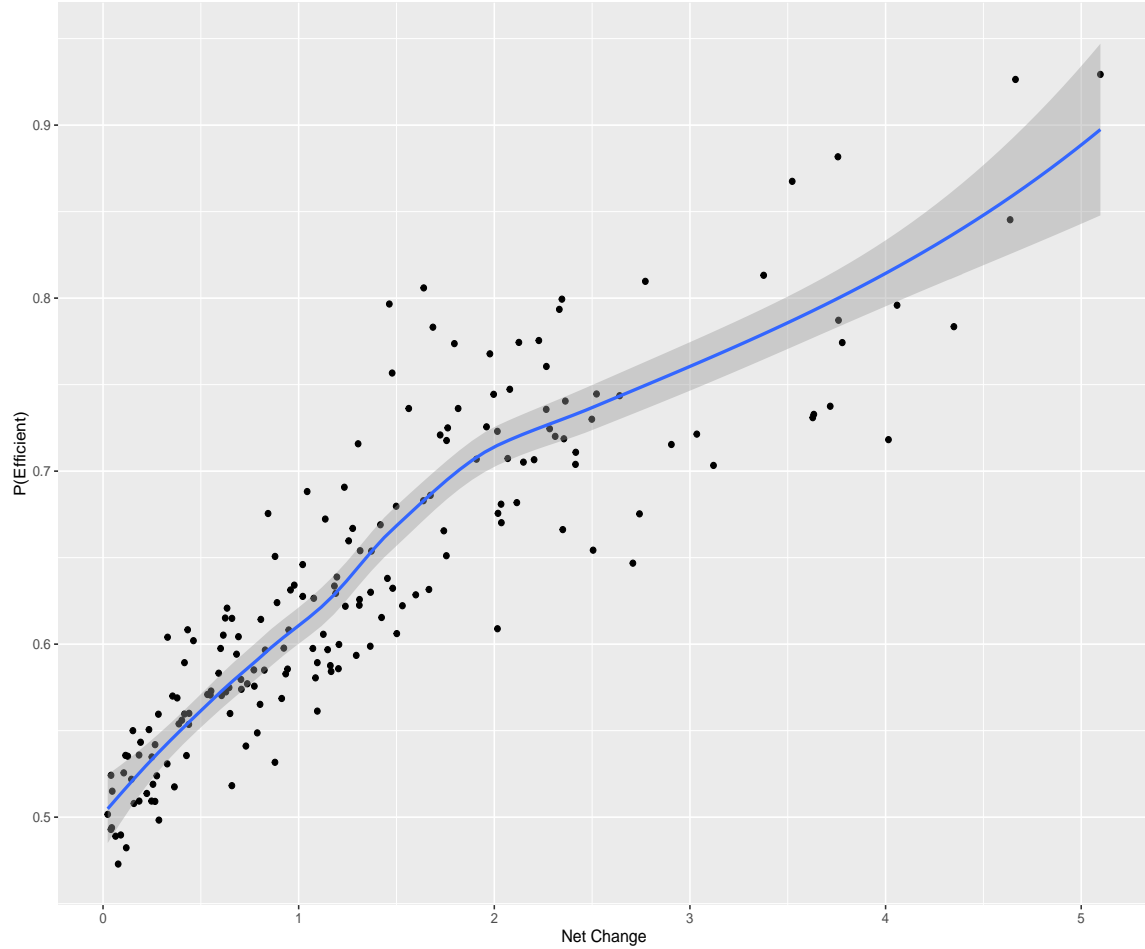
We see a clear relationship between the size of our original calculated net gain (without the random shocks) and the probability that the trade remains efficient when subject to the random shocks. From the plot we observe that the probability of an efficient trade begins around 47% for the smallest calculated net gains but increases linearly as the size of our net gains increase. The first and third quartiles of efficiency probabilities are calculated to be 57% and 71% respectively. Thus, 25% of our observed trades have at least at 71% probability of being efficient and 75% of our observed trades have at least at 57% efficiency probability when subject to random shocks. We also calculate the skewness of the efficiency probability

distribution to be 0.577, indicating the distribution is slightly skewed to the right. The full probability efficiency distribution is illustrated in the histogram below.

Figure 3.7: Histogram of P(Efficient)
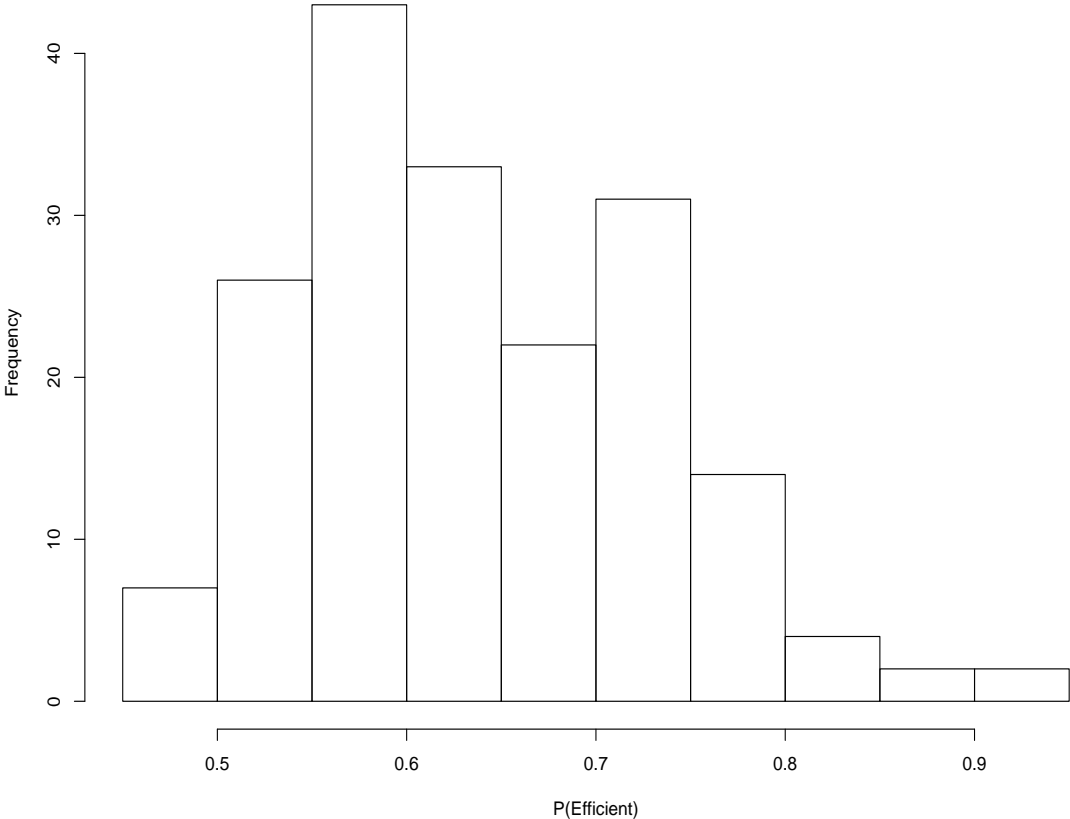
3.6 Conclusions

This paper uses novel approaches to measure team productivity and trading efficiency in Major League Baseball through an application of portfolio theory. Using a structural approach relating player outcomes to runs, we accurately capture the contributions of each

individual player to his team's output. We then build upon the applications of portfolio theory that have been implemented in other disciplines to define a team as a portfolio of assets that can be characterized by its mean and variance in production.

Using this characterization, we find that a large portion of the variation in team runs scored can be explained by our measures of mean and variance. It is observed that both mean and variance are positive and significant contributors to team output, with mean being a stronger determinant of team runs scored. The positive effect of variance arises from the sequential nature of production required to score runs in baseball, as higher variance teams tend to require less coordination between workers to produce output (runs) compared to lower variance teams.

Using this framework to characterize team productivity by team mean and variance, this paper makes use of a unique setting in professional baseball, where player contracts are considered intangible fixed assets and player productivity can be easily observed and objectively measured. Using this setting, we propose an approach to examine efficiency in the trade market for intangible fixed assets in baseball. A set of isorun curves for different eras over our timeframe are constructed to depict the mean-variance tradeoff needed to maintain a given level of team production in the form of runs per game. We then use this tradeoff to measure the trade efficiency for teams given their pre-trade expected values and variances and the traded player's pre-trade expected value and variance. It is observed that trade efficiency has improved over our timeframe and quite dramatically in the more recent years of our timeframe. This finding is consistent with the growth of analytics and data-driven decision making used in MLB front offices during this time.

In the areas of sport economics and sport finance, this framework could be extended to more continuous team sports such as basketball, soccer, football, hockey, etc. This approach could also be used to investigate more granular levels of performance by treating teams and players as portfolios of different skillsets. In baseball, an interesting application could look

at the success of MLB pitchers by characterizing pitchers by their portfolios of different pitch types. Another interesting extension could integrate player streakiness into the analysis to determine whether the specific sequence in which players achieve their successes over the course of the season have an impact on the overall levels of team production.

However, the implications of this study go beyond the professional sports industry as the framework developed in this paper can be extended to any team setting in which the performance of team members can be reliably measured. Of course, the impact of individual worker variance on team productivity will depend greatly on the specific process required for a team to produce output in a given industry and not all processes will require the amount of worker coordination needed to produce output as on a baseball team. The production process of a baseball team is one specific setting in which performance variance of individual team members has an overall positive impact on team productivity while controlling for the mean, due to the interdependence of individual worker outcomes needed to produce team output in baseball.

While there surely exist other industry-specific production processes for which worker variance has a lesser or negative effect on team production, our results illustrate the need for team managers to consider the impact that individual worker variance has on the overall levels of team production. In addition, these results suggest that managers should account for the variance in worker productivity in addition to mean levels of production when completing performance evaluations for workers. Our findings indicate that worker variance is particularly relevant in any industry in which sequential production occurs and a high degree of worker coordination is needed.

Chapter IV

Entry Level Job Assignments and Career Length: The Case of Major League Baseball Pitchers

4.1 Introduction

The impact of early career decisions and opportunities on future career paths for workers has been widely covered across disciplines in the academic literature. Most previous research has focused on examining the impact of variables such as labor market conditions, employee attributes, and firm characteristics on an individual's future earnings and opportunities throughout the worker's career. One related area of research that has yet to be investigated thoroughly is the impact of employer-assigned roles on a worker's future career opportunities.

Altonji (2005) is one of the few studies to investigate this question by theoretically showing how statistical discrimination by employers can have a detrimental effect on the career paths of workers with limited initial information available to the employer. These workers are often placed in lower-skilled roles initially, denying any high-skilled workers an opportunity to quickly establish themselves as such. Salaga et al. (2020) also conduct a study looking at the impact of employer role assignments on future career outcomes in a sport context. The authors focus on determining the impact of employer-assigned workload on future

career length and productivity for running backs in the National Football League (NFL). It is found that high workload levels assigned to players by teams have a negative impact on the player's future productivity, but a positive effect on the player's career length.

This paper seeks to build upon these studies by investigating the impact of employer-assigned, entry-level roles on the career paths and survivability of pitchers in Major League Baseball (MLB). The pitcher position in MLB provides us with several desirable characteristics of conducting such a study. First, there exist two distinct primary pitcher roles in professional baseball, starting pitchers and relief pitchers, where starting pitchers have historically been thought to be more valuable.[1] Second, significant differences in the labor supply exist for the two types of pitcher, with relief pitchers having a much larger supply of labor. Third, a pitcher's entry-level role is largely determined by the team's subjective evaluation of the best role for the pitcher on the team given the pitcher's skillset and existing team composition. Finally, transitions between roles are common for pitchers. As a result, this allows us to better investigate and isolate the true impact of a pitcher's entry-level role on his career path and survivability.

We hypothesize that an entry-level relief pitcher assignment will have a detrimental impact on a pitcher's survivability in the league as the large labor supply of relief pitchers should make it easier for teams to replace relief pitchers with affordable alternatives compared to starting pitchers. Additionally, given that starting pitchers have historically been viewed as the more valuable type of pitcher, an entry-level relief pitcher assignment may lead to a perception of lower-ability for the pitcher by potential future employers.

Kaplan-Meier estimates are calculated for entry-level starting pitchers and entry-level relief pitchers and a log-rank test for the equality of survivor functions is conducted to determine if significant survival differences exist by entry-level pitcher type. Additionally, a Cox-proportional hazard model is run to allow for multivariate modeling. We find that

---

[1]https://https://www.baseball-reference.com/bullpen/Relief_pitcher

entry-level relief pitchers face a significant survival disadvantage compared to entry-level starting pitchers and that this disadvantage essentially persists for the duration of an entry-level relief pitcher's career. We also find that significant survival disadvantages (advantages) continue to exist for entry-level relief (starting) pitchers who transitioned roles to starting (relief) pitchers early in their careers. The results of this study suggest that the early-career decisions made by teams on behalf of their players ultimately has career altering implications for these players.

This paper proceeds as follows: Section 2 reviews the previous literature related to this topic and presents further background information on our pitcher types. Section 3 develops the approach to be used in this paper and summarizes our results. Section 4 provides a brief discussion with our conclusions.

4.2   Background and Previous Research

The impact of early career decisions and opportunities on employee career paths has been covered extensively in the academic literature across industries. Several of these studies focus on how early career labor market conditions can have an impact on a worker's career outcomes. Gebel (2003) investigates the impact of early-career temporary employment on subsequent career wages and temporary employment cycles. The author finds short-term wage and employment penalties for those who enter temporary employment positions early in their careers. Raaum and Røed (2006) show that the labor market conditions at the time of job-market entry have a long-term impact on an individual's future employment opportunities. Von Wachter and Bender (2006) study the effects of early career job loss on worker earnings and find evidence of short-term losses for workers and long-term losses for those workers who were forced to transition from large firms to smaller organizations. Ahlin et al. (2014) determine the impact that market thickness for skills has on the initial

71

wages of recent graduates and find that graduates are more likely to move to urban areas for employment after graduating where thicker markets for skills exist. Schmillen and Umkehrer (2017) focus on the impact of early-career unemployment on future unemployment and find that individuals that experience unemployment early in their careers are more likely to experience future unemployment as well.

Other studies focus on the impact of more specific firm or employee characteristics on an individual's career path. Kelley (1978) shows that individuals coming from wealthy families tend to have advantages in early career employment compared to individuals from lower-class families. The author also shows that these advantages persist for an individual throughout their careers. Mattox and Jinkerson (2005) investigate the impact of company-provided training for new hires on the retention of experienced workers. Van der Heijden et al. (2009) study the impact of age effects on career success while Stumpf et al. (2010) look at how the practices of human resources impact employee career success. Biemann and Braakmann (2013) determine the impact of international career experience on early career success while Dossinger et al. (2019) look at the impact of physical attractiveness of early-career salaries. Musset and Kurekova (2018) illustrate the importance of career guidance in the early stages of an individual's education. Monti et al. (2020) investigate the impact of previous employer characteristics on the gender-wage gap and find that a significant portion of the wage gap can be attributed to women being overrepresented in lower-paying firms compared to men.

A variety of industry specific studies on early-career outcomes and career paths have been conducted as well. Cox and Harquail (1991) find that starting salaries and starting job levels of MBA graduates have a significant impact on measures of career success. Dougerty et al. (1993) show that MBA graduates are less likely to change jobs early in their careers than bachelor's degree graduates. Williamson and Cable (2003) investigate the determinants of early-career research productivity of management faculty. Quartz et al. (2008) conduct a study of early-career urban educator retention and role changing while Cochran-Smith et al.

(2012) investigate how early-career experiences influence a teacher's decision to stay or leave the discipline. DeAngelis et al. (2013) consider how the support and preservice preparation of early career teachers influence these decisions as well. Goldacre et al. (2010) compare the early career specialty decisions of recent medical graduates with their eventual specialty decisions. Flinkman et al. (2013) investigate the factors that lead to early-career nurses leaving the profession.

However, one limitation of this previous literature thus far is that objective measures of employee productivity are difficult to observe and control for. As a result, several sport economics studies pertaining to worker survivability and career paths have been conducted. Labor market studies are prevalent in sport economics for several reasons as outlined by Kahn (2000). In short, the professional sports industry provides researchers with several desirable characteristics for conducting labor market research that are often difficult to find in most other industries. The characteristics include observability of worker performance, precise measures to quantify worker productivity, and public information on worker wage data.

Groothuis and Hill (2004) find that player performance metrics are a significant determinant of career length in the National Basketball Association. Additionally, the authors do not find any evidence of exit discrimination by race in the league, contrary to the findings of Hoang and Rascher (1999). Witnauer et al. (2007) look at career lengths of MLB players in the 20th century and observe that the career length distribution for players is skewed toward shorter career lengths, while Petersen et al. (2011) make similar observations pertaining to short careers in MLB and the NBA. Baker et al. (2013) investigates whether career length differences exist by player position for different sports. The authors find that career length differences by position were present in MLB and the National Football League (NFL), but not the NBA and National Hockey League (NHL). Ducking and Bollinger (2014) determine whether minimum salaries have an impact on career length in the NFL and find that players

making the league minimum have significantly shorter careers.

Despite the abundance of literature on early career opportunities, career paths, and survivability across a variety of disciplines, one area that has yet to be researched extensively is the impact of employer-assigned job functions on an employee's career length. As previously mentioned, Altonji (2005) and Salaga et al. (2020) are two existing studies that address issues relevant to this area.

This paper seeks to build upon these studies investigating the impact of employer-assigned roles on worker career outcomes by utilizing a scenario from baseball with desirable characteristics for conducting such a study. Specifically, our goal is to determine whether a pitcher's entry-level role, as determined by his team, has an impact on the pitcher's ability to survive in the league. In professional baseball, pitchers are typically categorized into one of two types of pitching positions: starting pitcher or relief pitcher. Starting pitchers start the game as his team's pitcher and generally pitch the majority of innings in the game, while relief pitchers generally "relieve" the starting pitcher at the manager's discretion due to either to fatigue or ineffectiveness of the starting pitcher. Additionally, role transitions between starting pitchers and relief pitchers are common as approximately 20% of pitchers transition from a primary starting pitcher to a primary relief pitcher, or vice versa, within the first three seasons of their careers.

Over the timeframe 1976-2016, the average length of an appearance for a starting pitcher was approximately 6.04 innings while the average length of an appearance for a relief pitcher during this timeframe was approximately 1.23 innings. As can be inferred from these averages, generally a combination of relief pitchers are used to finish a game once the starting pitcher has departed. Additionally, given that starting pitchers tend to pitcher more innings over the course of a season than a relief pitcher, starting pitchers have generally been thought of as the more valuable type of pitcher and typically earn higher salaries than relief pitchers over our timeframe.

However, as Schmidt (2020) alludes to, there have been drastic changes to the role of relief pitchers in Major League Baseball over the past couple of decades. The following two charts illustrate this trend by respectively showing the percentage of MLB innings pitched by relief pitchers and the number of relief pitchers that appeared in an MLB game during the season for 1976-2016.

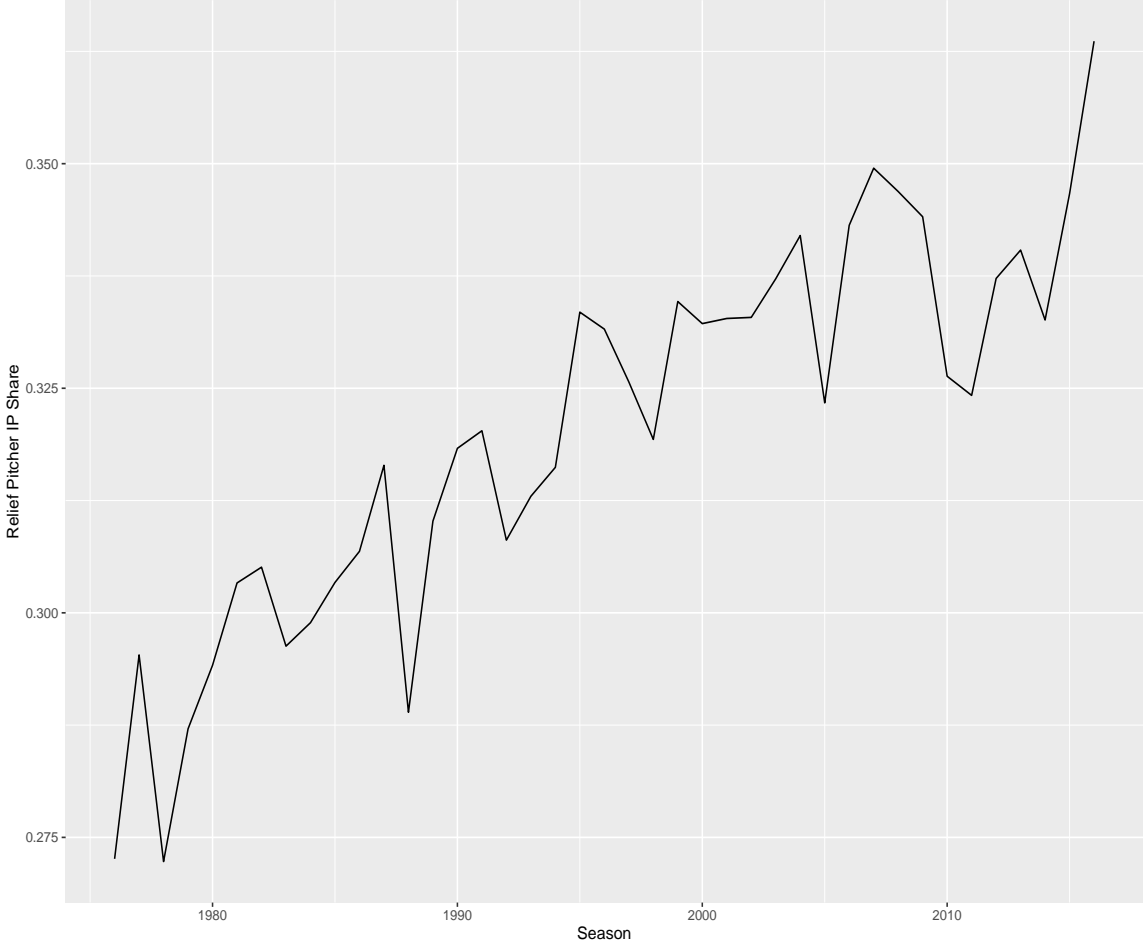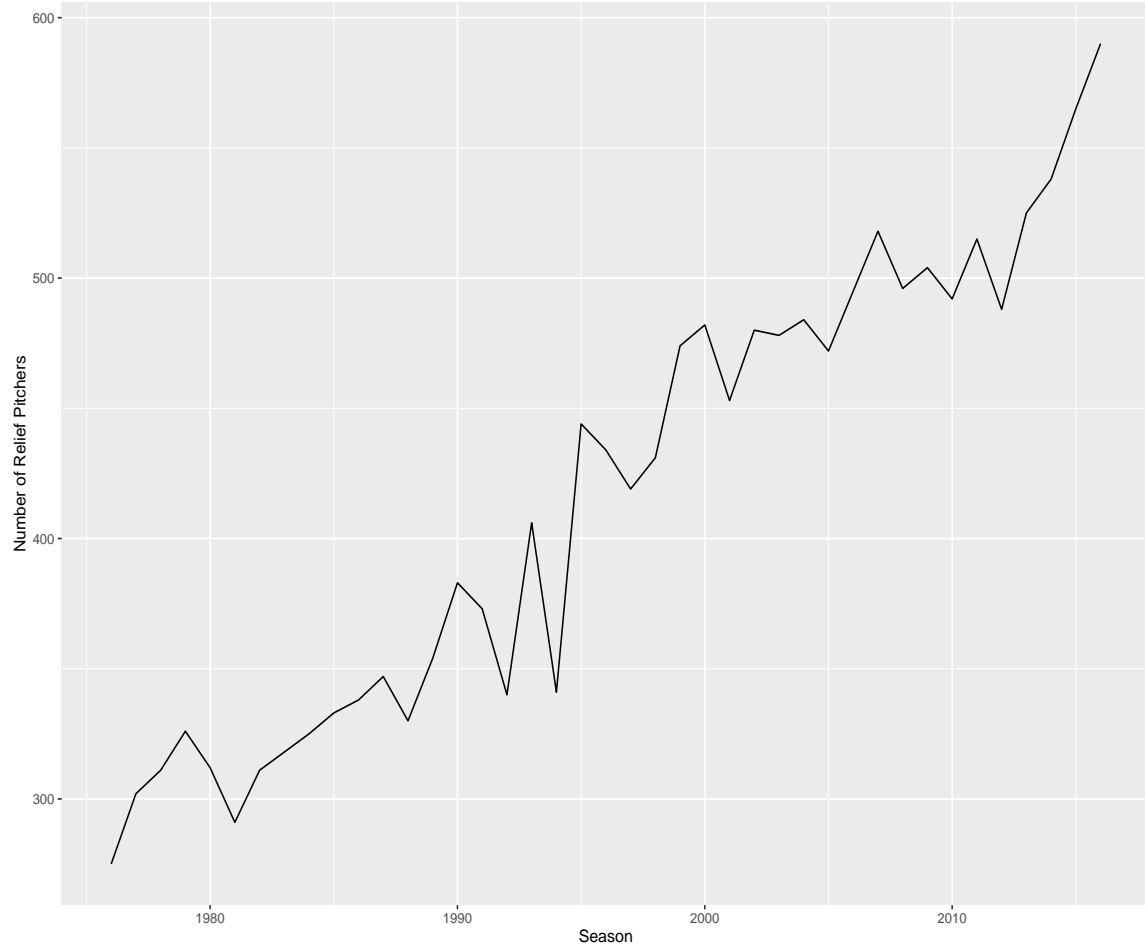Figure 4.1: Relief Pitcher Share of Innings Pitched by Season: 1976-2016

Figure 4.2: Number of Relief Pitchers by Season: 1976-2016

Over this timeframe, the share of total innings pitched by relief pitchers has increased by about nine percentage points, while the aggregate total number of innings pitched from 9521.1 in 1976 to 15747.6 in 2016, an increase of 92.47% (though some of the increase in aggregate innings is a result of league expansion over this time). Over this same period, the number of relief pitchers that made at least one appearance in the league during the season increased from 275 in 1976 to 590 in 2016, an increase of 114.55%. For reference, the same calculation for starting pitchers, infielders (first basemen, second basemen, third basemen, and shortstops) and outfielders resulted in increases of 37.25%, 28.83%, and 41.81%

respectively, while the number of players that made at least one appearance at each position type in 2016 were 280 for starting pitchers, 362 for infielders, and 329 for outfielders. Clearly the labor supply of relief pitchers far exceeds that of any other position in professional baseball over this timeframe.

Furthermore, over this timeframe roster sizes remained constant at 25 while five-men starting pitcher rotations were the norm.[2] The number of relief pitchers teams typically carried on their roster at any given time increased only slightly from around 5 to 8.[3] Thus, given the relatively minor changes to roster sizes over this timeframe, it can be implied that the amount of roster turnover for relief pitchers has increased substantially over this period. Intuitively this makes sense, since if a surplus of talent exists due to the large labor supply of relief pitchers, then the average relief pitcher should be much easier to replace than essentially any other position over our timeframe. This is consistent with diminishing productivity differentials (i.e. talent compression) as discussed in Gould (1986), Chatterjee and Yilmaz (1991), Butler (1995), Schmidt and Berri (2003), and Schmidt and Berri (2005), Zimbalist (2010), and Schmidt (2020), due to increases in the labor supply.

Given the previous academic research on the implications of early career opportunities and labor market conditions on a worker's career path, we hypothesize that an entry-level relief pitcher assignment will negatively impact a player's ability to survive in the league compared to an entry-level starting pitcher assignment. Specifically, we seek to address four primary questions pertaining to this hypothesis: (1) Do entry-level relief pitchers have shorter careers on average than entry-level starting pitchers? (2) How long do any survival differences in persist? (3) Do early-career role transitions affect survival differences between entry-level pitcher types? (4) Are differences in survival for entry-level pitcher types explained by differences in ability? Our empirical approach and results pertaining to these questions are

---

[2]https://www.baseball-reference.com/bullpen/Pitching_rotation
[3]https://www.baseball-reference.com/bullpen/Roster

included in the following section.

4.3   Approach and Estimation

To address our research questions we utilize the Kaplan-Meier estimator of the survival function for our pitcher classifications based on their entry-level roles in their debut seasons. We obtain our player debut and career length data from the Lahman Database for the seasons 1976-2016.[4] For every pitcher that made his MLB debut over this timeframe, our data includes the pitcher's entry-level role and career length. We classify a pitcher as an entry-level relief pitcher if he makes at least 50% of his appearances as a relief pitcher in his debut season. Otherwise, we classify the pitcher as an entry-level starting pitcher. After censoring for active players, our sample includes 3,702 total pitchers over this timeframe.

The Kaplan-Meier estimator for the survival function is calculated as follows:

$$(4.1) \qquad \hat{S}(t) = \prod_{i:\, t_i \,\leq\, t} (1 - \frac{d_i}{n_i}),$$

where t_i represents time in league (career length), d_i represents the number of players that do not survive in the league past career length t_i, and n_i represents the number of players that have survived to career length t_i.

Based on our Kaplan-Meier estimates for each classification of pitcher, we then construct a set of Kaplan-Meier curves to display the survival rate for each group over each time increment (career length in years) in our data. After constructing our Kaplan-Meier curves, we test to determine if significant differences exist in the survivor functions of entry-level starting pitchers and entry-level relief pitchers using a log-rank test with a null hypothesis of equal survivor functions between groups, where the test statistic follows a chi-square

---

[4]http://www.seanlahman.com/baseball-archive/statistics/

distribution. Using this framework, we address our first three primary research questions.

4.3.1 Do significant differences exist in survivor functions between entry-level starting pitchers and entry-level relief pitchers?

The Kaplan-Meier curves for entry-level starting pitchers versus entry-level relief pitchers are included below.

Figure 4.3: Kaplan Meier Curves: Entry-Level Starting Pitchers vs. Entry-Level Relief Pitchers (1976-2016)
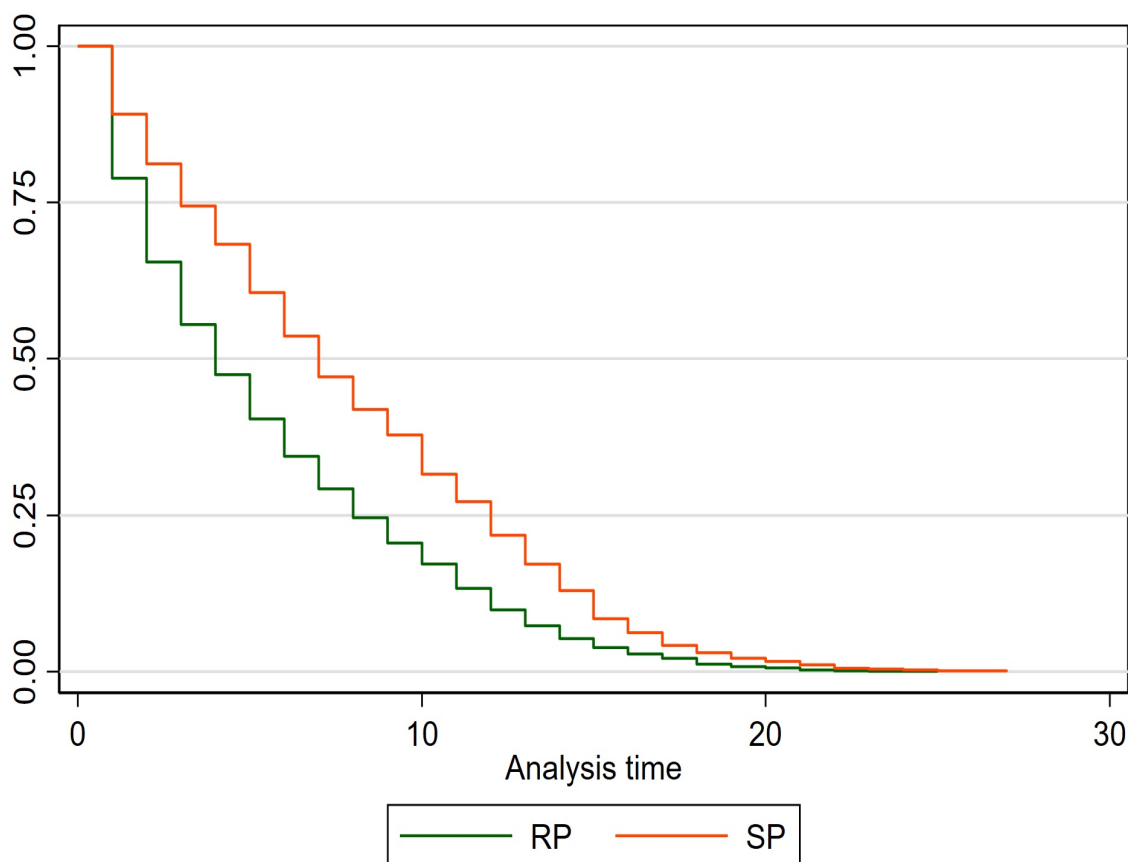
Our log-rank chi-square test statistic is equal to 159.16, which is significant at the .001 level. These results confirm that significant differences exist in survivor functions based on

a pitcher's entry-level role in the league. We see a significantly weaker survival rate and shorter career length for entry-level relief pitchers compared to what we observe with entry-level starting pitchers. In particular, we see that the survival disadvantage for entry-level relief pitchers is striking in the early years of one's career. This result certainly aligns with our hypothesis of entry-level relief pitchers facing a tougher survival curve due to the large labor supply and ease of replaceability. Of course, a potential endogenous variable with these results is the overall pitcher quality of entry-level starting pitchers versus entry-level relief pitchers which we will address later.

4.3.2 How long do differences in survivor functions persist?

Given the striking disparity between the survival rates of entry-level relief pitchers and entry-level starting pitchers at the early stages of their careers, we would now like to determine how long significant differences in survival functions exist between entry-level pitcher types when conditioning on career length. In particular, we are interested in determining whether the differences in survival functions can be attributed entirely to reliever attrition in the early stages of their careers, or whether these survival disadvantages for entry-level relief pitchers persist throughout these pitchers' careers. Stated another way, we would like to determine if an experience threshold exists for which the survival curves between entry-level relievers and entry-level starting pitchers are similar thereafter.

To do this, we calculate the Kaplan-Meier estimates for entry-level relief pitchers and entry-level starting pitchers once again while now conditioning our data on players that have survived at least X years in the league. We complete our log-rank test for every value of $X = 2, \ldots, 15$ to determine where the experience threshold for similar survival functions exists for entry-level relievers and starting pitchers. We choose 15 years of experience as an endpoint as this represents the 95th percentile of pitcher career length in our data. We have included our log-rank chi-squared test statistics below.

Table 4.1: Log-Rank Test Results Conditional on Career Length

X	Log-Rank Test Statistic	N (SP)	N (RP)	N
2	104.48***	1032	1956	2988
3	70.80***	928	1593	2521
4	47.33***	840	1322	2162
5	30.65***	760	1105	1865
6	24.75***	659	913	1572
7	21.05***	568	754	1322
8	18.35***	486	619	1105
9	12.87***	423	503	926
10	6.14*	375	406	781
11	7.73**	305	329	634
12	2.34	257	244	501
13	0.58	201	173	374
14	0.04	154	123	277
15	0.04	113	85	198

Test Statistic Follows Chi-Square Distribution

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

From the results we observe that entry-level relief pitchers and entry-level starting pitchers do not face similar survival curves until they reach 12 years in the league. While we see that the gap in survival rate between the two groups does tend to shrink with league tenure, our log-rank test finds that statistically significant survival differences exist for essentially the duration of a pitcher's career (a career length of 12 years represents the 92nd percentile of entry-level relief pitcher career lengths in our data). This result implies that a team's decision to assign a pitcher a primary relief pitcher role in his debut season has career-altering implications for the player's survivability in the league, as the majority of entry-level relief pitchers will face a survival disadvantage compared to entry-level starting pitchers that persists throughout their careers.

4.3.3 How do early career role transitions affect the survival differences between pitcher types?

Thus far we have established that pitchers assigned to entry-level reliever roles face a much tougher survival curve than entry-level starting pitchers, and that a statistically significant survival disadvantage for relief pitchers essentially lasts for the duration of their careers. However, one additional factor we have yet to consider is the impact of early-career role transitions on career length. Post-debut role transitions between starting pitchers and relief pitchers are fairly common as approximately 20% of the pitchers in our data transitioned their primary roles within the first three seasons of their careers. This provides us with an opportunity to determine whether early career role transitions from relief pitcher to starting pitcher eliminate the survival disadvantage entry-level relief pitchers face. Similarly, we can determine whether early career role transitions from starting pitcher to relief pitcher eliminate the survival advantage entry-level starting pitchers typically have. If significant survival disadvantages (advantages) continue to exist for entry-level relief (starting) pitchers after an early career transition to a starting (relief) pitcher role, this would support the notion of teams perceiving entry-level relievers to be of lower ability than entry-level starting pitchers. This would also indicate that differences in survival cannot be explained entirely by differences in labor supply.

To determine whether early career role transitions from relief pitcher to starting pitcher eliminate the survival disadvantage for entry-level relief pitchers, we estimate and compare the Kaplan-Meier survival functions of entry-level relief pitchers who transitioned to starting pitchers within the first three years of their careers to entry-level starting pitchers who did not transition out of a starting pitcher role within the first three years of their careers. Similarly, to determine whether early career role transitions from starting pitcher to relief pitcher eliminate the survival advantage for entry-level starting pitchers, we estimate the

Kaplan-Meier survival functions of entry-level starting pitchers who transitioned to relief pitchers within the first three years of their careers to entry-level relief pitchers who did not transition out of a relief pitcher role within the first three years of their careers. In both cases we only include pitchers with a career length greater than one year in our data since pitchers that transitioned roles must have a career length of at least two under our specification. For both comparisons, we again conduct a log-rank test for the equality of survivor functions. We begin with our comparison of early career entry-level reliever transitions to entry-level starting pitchers.
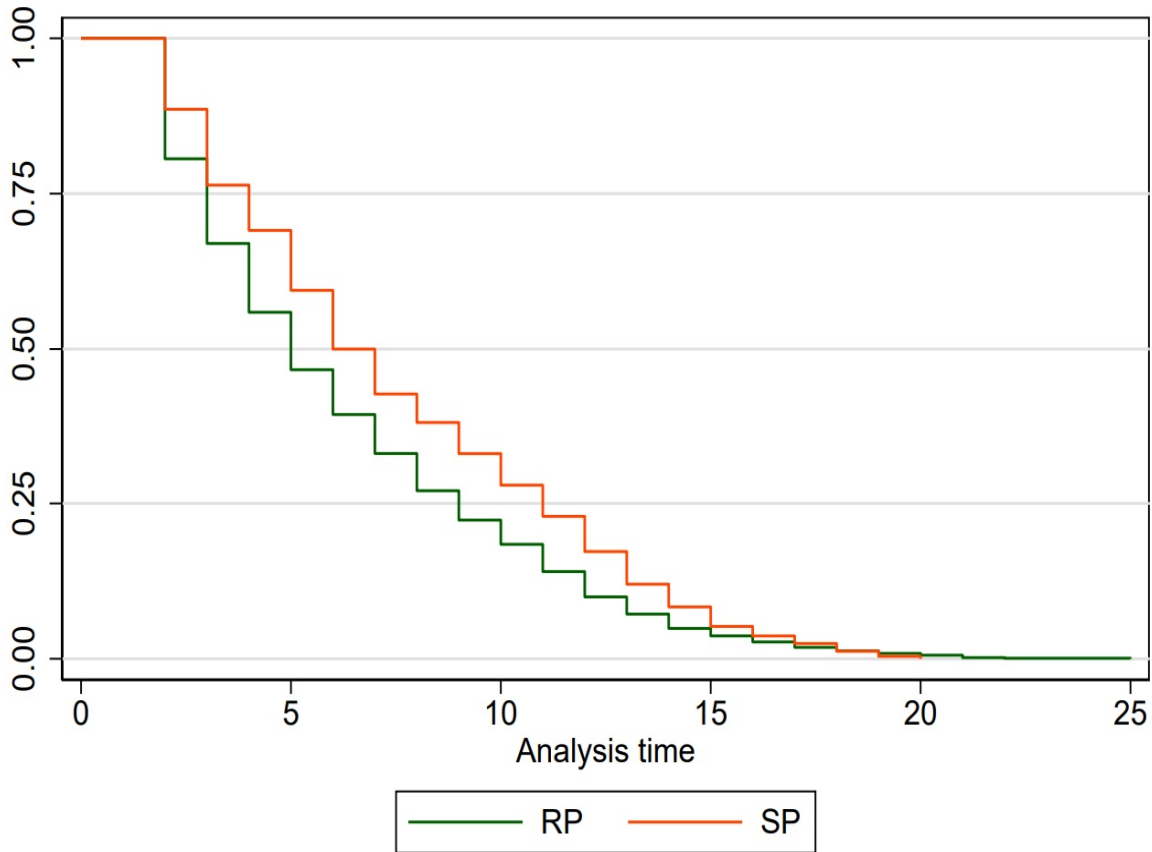
Figure 4.4: Kaplan Meier Curves: Entry-Level Relief Pitcher to Starting Pitcher Transitions vs. Entry-Level Starting Pitchers (1976-2016)
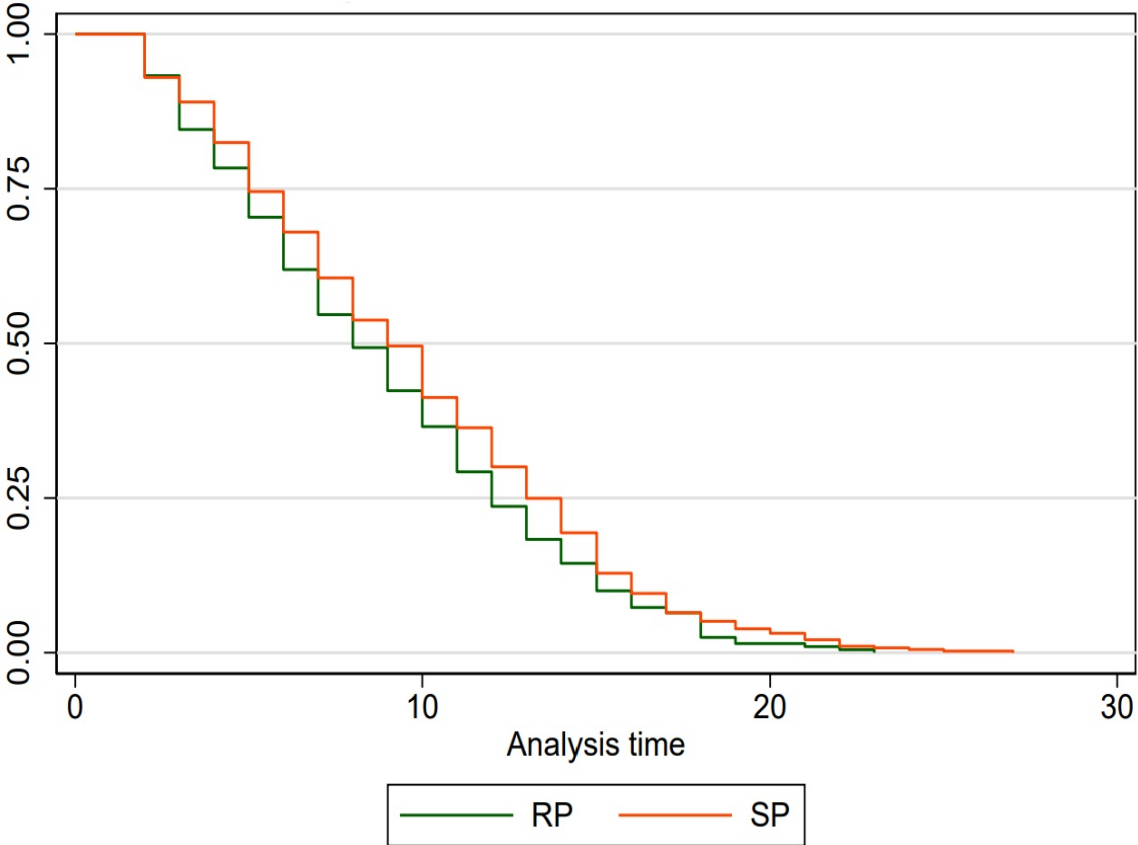
Our chi-square test statistic for our log-rank test is equal to a value of 5.94, significant at the .05 level. In our comparison of early career entry-level reliever to starter transitions and entry-level starting pitchers, we do observe that significant differences in survival estimates persist, with entry-level relief pitchers continuing to experience a significantly shorter career length. However, it should be noted that the survival disadvantage between entry-level relievers that transitioned roles early in their careers and entry-level starting pitchers has decreased in relation to what we observed in the results comparing all entry-level relievers to all entry-level starting pitchers, indicating that the early career transition to from relief pitcher to starting pitcher does positively affect the pitcher's career length. Nevertheless, the fact that statistically significant disadvantages in survival exist for entry-level relief pitchers despite transitioning to starting pitchers early in their careers is consistent with what we would expect to see if teams viewed an entry-level reliever assignment as a signal of lower ability.

Now we turn our attention to the Kaplan-Meier curves comparing early career entry-level starting pitcher transitions to entry-level relief pitchers.

Figure 4.5: Kaplan Meier Curves: Entry-Level Starting Pitcher to Relief Pitcher Transitions vs. Entry-Level Relief Pitchers (1976-2016)

The chi-square test statistic for our log-rank test is equal to 24.45, significant at the .001 level. The results comparing early career entry-level starting pitcher to relief pitcher transitions to entry-level relief pitchers support our previous results of survival advantages for entry-level starting pitchers. We do observe that the survival advantage between entry-level starters that transitioned to relievers early in their careers and entry-level relief pitchers has decreased in relation to what we observed in the results comparing all entry-level relievers to all entry-level starting pitchers. This suggests that the early career transition to from starting pitcher to relief pitcher has a negative impact on the pitcher's career length. However, entry-level starting pitchers continue to have a statistically significant survival advantage

compared to entry-level relief pitchers despite the fact that these entry-level starting pitchers transitioned to relief pitcher roles early in their careers.

These are striking results that illustrate the importance of a pitcher's entry-level role on his career path despite role transitions that might occur in the early stages of his career. Given the significance of these findings we decide to conduct our Kaplan-Meier transition analysis for a couple alternate starting pitcher and relief pitcher classification definitions as a check for robustness. The first alternate definition classifies a relief pitcher as a pitcher with at least 25% of his appearances as a reliever and a starter otherwise. This approach makes it more difficult to be classified as a starting pitcher and easier to be classified as a relief pitcher. The second alternate definition classifies a relief pitcher as a pitcher with at least 75% of his appearances as a reliever and a starter otherwise. This approach makes it more difficult to be classified as a relief pitcher and easier to be classified as a starting pitcher. In both cases a pitcher transition will be respectively determined using these alternate classification thresholds.

The Kaplan-Meier curves for the comparison of early career entry-level reliever transitions to entry-level starting pitchers for both alternate definitions are included below.

Figure 4.6: Kaplan Meier Curves: Entry-Level Relief Pitcher (at least 25% of appearances) to Starting Pitcher (at least 75% of appearances) Transitions vs. Entry-Level Starting Pitchers (1976-2016)
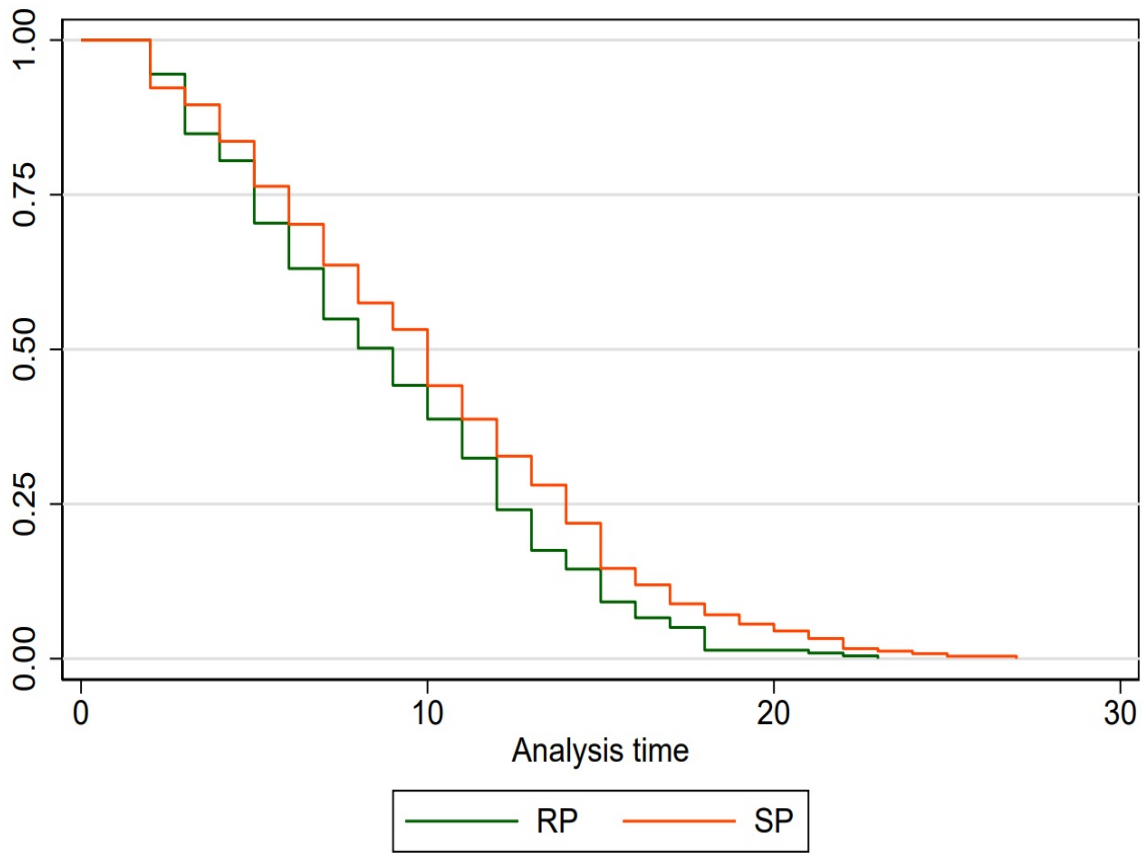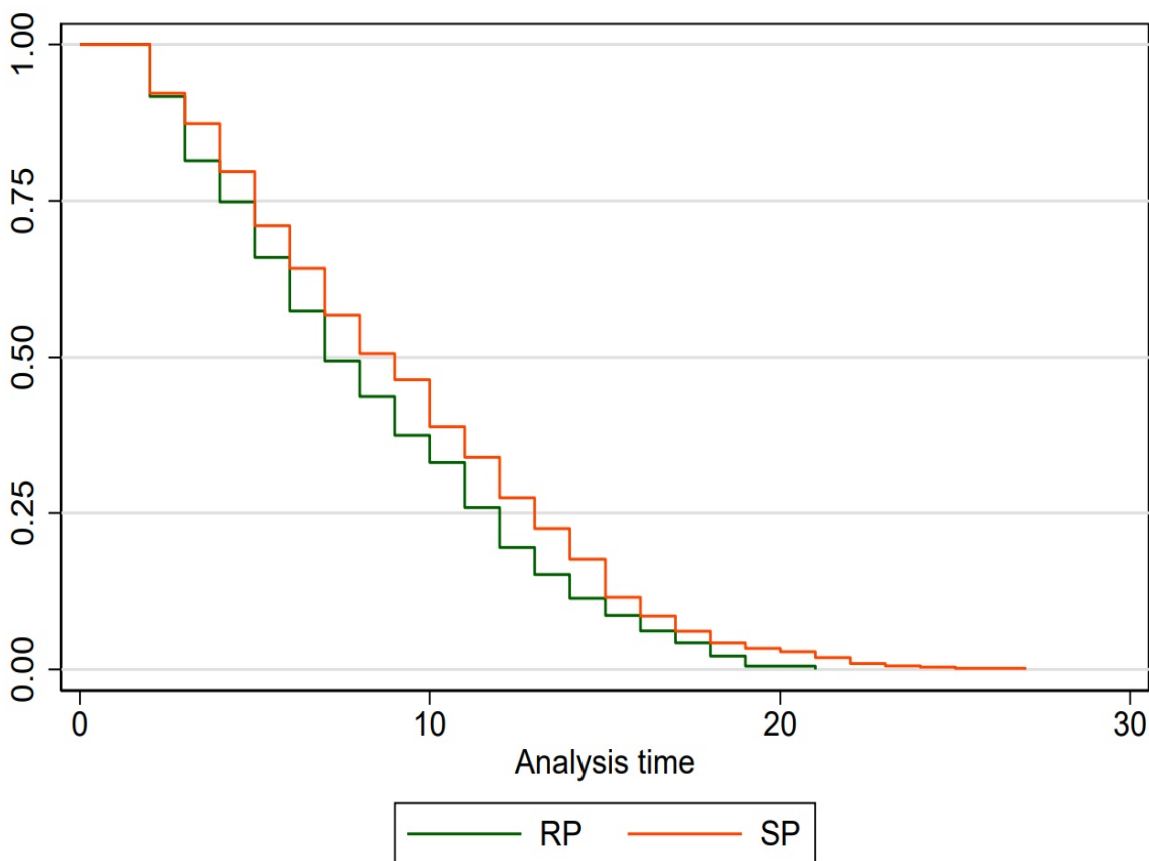
Figure 4.7: Kaplan Meier Curves: Entry-Level Relief Pitcher (at least 75% of appearances) to Starting Pitcher (at least 25% of appearances) Transitions vs. Entry-Level Starting Pitchers (1976-2016)

The chi-squared statistics for the log rank tests pertaining to the Kaplan-Meier curves for Figures 4.6 and 4.7 respectively are 10.88 and 10.19, both significant at the .01 level. Thus, our previous result of entry-level relief pitchers continuing to face a survival disadvantage despite transitioning to starting pitchers early in their careers is robust to these alternate pitcher classification definitions. We now test whether our results for the comparison of early career entry-level starting pitcher transitions to entry-level relief pitchers are robust to these alternate definitions. The corresponding Kaplan-Meier curves are included below.

Figure 4.8: Kaplan Meier Curves: Entry-Level Starting Pitcher (at least 25% of appearances) to Relief Pitcher (at least 75% of appearances) Transitions vs. Entry-Level Relief Pitchers (1976-2016)
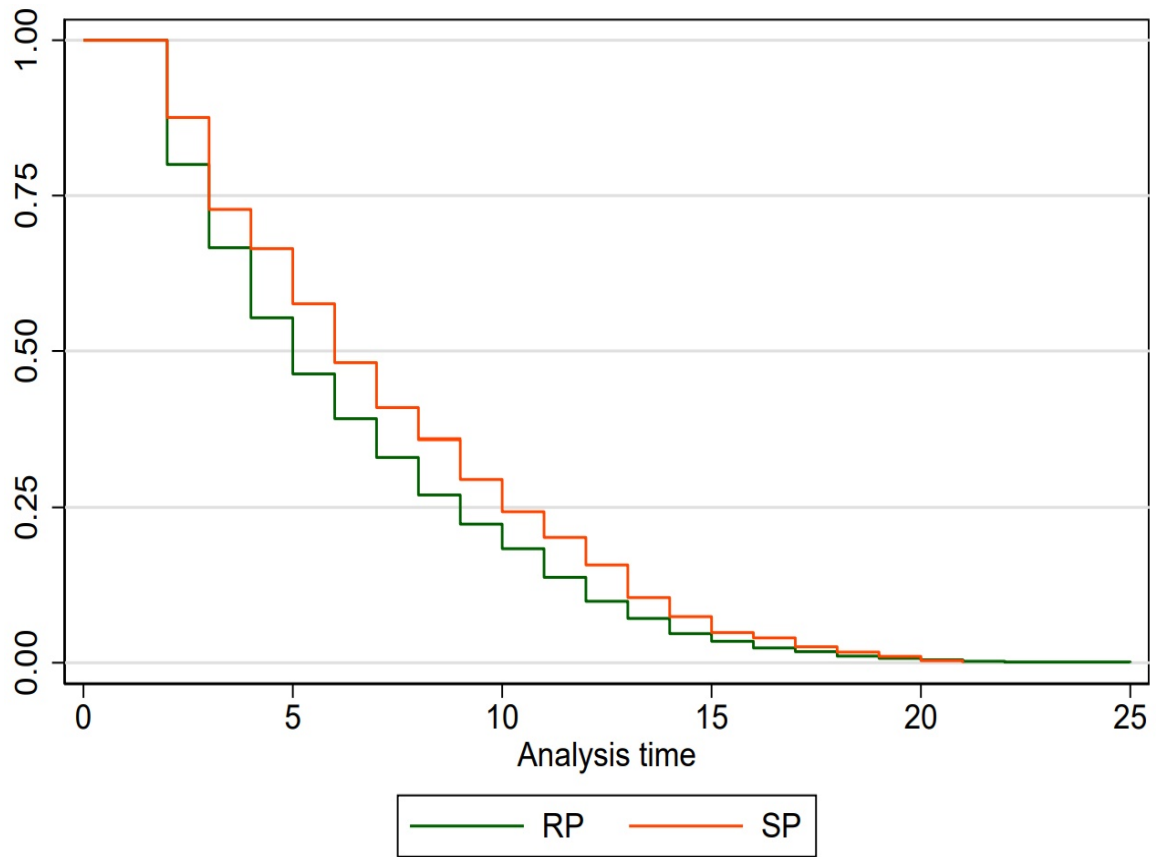
Figure 4.9: Kaplan Meier Curves: Entry-Level Starting Pitcher (at least 75% of appearances) to Relief Pitcher (at least 25% of appearances) Transitions vs. Entry-Level Relief Pitchers (1976-2016)
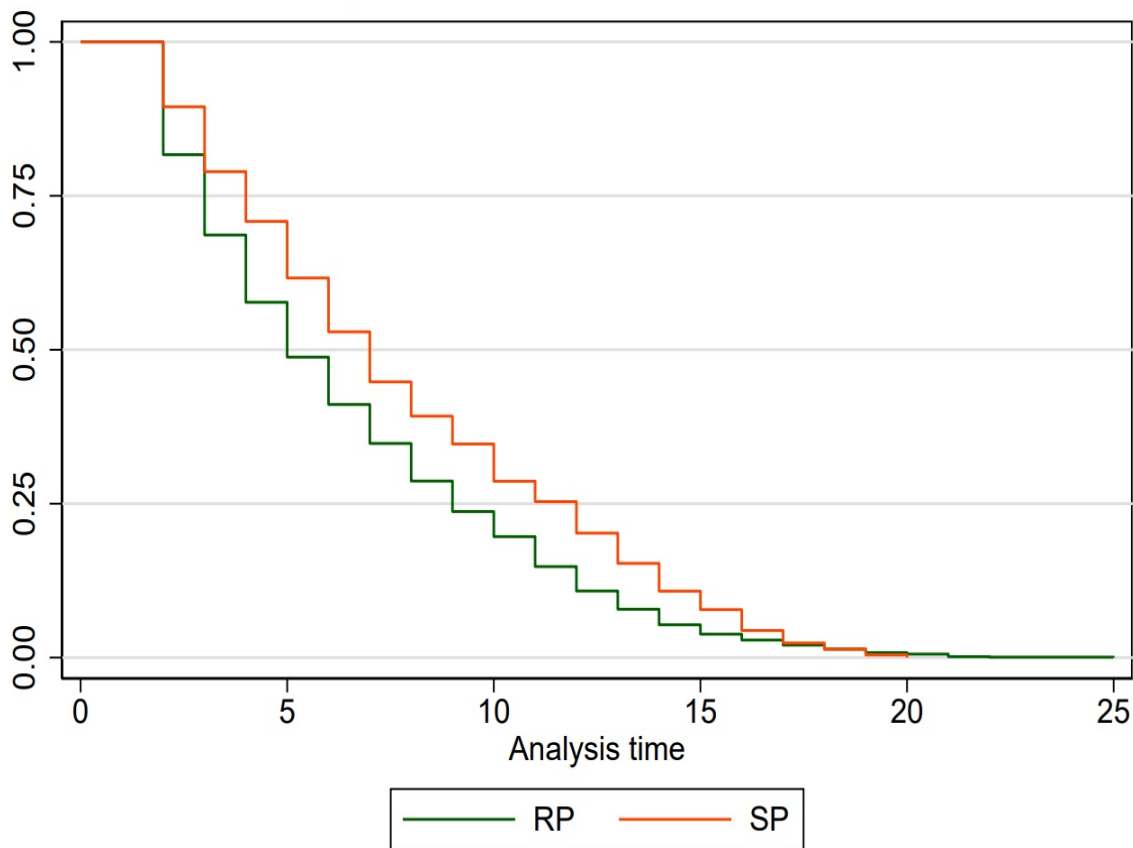
The chi-squared statistics for the log rank tests pertaining to the Kaplan-Meier curves for Figures 4.8 and 4.9 respectively are 16.06 and 26.81, both significant at the .001 level. Thus, our previous result of entry-level starting pitchers continuing to experience a survival advantage despite transitioning to relief pitchers early in their careers is robust to these alternate pitcher classification definitions.

4.3.4   Examining Differences in Pitcher Quality

While our previous results present evidence that an entry-level relief pitcher assignment is detrimental to a pitcher's ability to survive in the league compared to an entry-level starting pitcher assignment, we have yet to account for the impact of pitcher quality on these results. An endogeneity issue may exist with our current results if less talented, lower quality pitchers are systematically classified as relief pitchers by teams. If this were the case, then it could be reasonably argued that our results thus far are due primarily to differences in pitcher quality rather than the differences in labor supply or perceived ability for entry-level relievers that we have previously mentioned.

One reason for which we have not accounted for pitcher quality thus far is due to the structurally different role of starting pitchers and relief pitchers (Krautmann et al., 2003). Stamina tends to be valued much more for starting pitchers compared to relief pitchers as the average number of innings pitched for starting pitchers is more than twice that of relief pitchers during their entry-level seasons in our data. For this reason, a trade-off typically exists between innings pitched and other normalized rate metrics such as strikeout rate or earned run average for starting pitchers. Since stamina is not nearly as valued for relief pitchers, these normalized rate metrics for relief pitchers tend to be more highly valued by teams. Thus, a direct comparison of performance metrics for starting pitchers and relief pitchers can lead to misleading results pertaining to differences in quality by pitcher type.

However, we can make a direct performance comparison when focusing on pitcher transitions by comparing the post-transition season-level performance for entry-level relief pitchers who transitioned roles early in their careers to the season-level performance of entry-level starters who never transition roles in our data. We can analogously compare the post-transition season-level performance of entry-level starting pitchers who transitioned roles early in their careers to the season-level performance of entry-level relievers who never tran-

sition roles in our data. Once again we define an early career transition to be a role transition that takes place within the first three seasons of a pitcher's career. To compare pitcher quality in each case, we will run an independent two sample t-test to determine if significant differences exist between the two pitcher groups for season-level hit rates, strike out rates, walk rates, runs allowed per 9 innings, and innings pitched. Our data includes every pitcher season for these groups for which the pitcher had at least 5 innings pitched. This minimum number of innings pitched is used for the stability of our rate metrics. If systematic differences in pitcher quality between entry-level starting pitchers and entry-level relief pitchers exist, then we would expect to see that entry-level starting pitchers perform better than entry-level relief pitchers across performance metrics for both comparison sets. The corresponding t-test results are included below.

Table 4.2: Independent Two-Sample t-test: Transition Group - Early Career Entry-Level Relief Pitcher to Starting Pitcher Transitions

Variable	Control Group Mean	Transition Group Mean	Difference	Standard Error	t-value
Hit Rate	1.030	1.051	-0.0209***	0.00603	-3.464
Strikeout Rate	0.740	0.676	0.0644***	0.00630	10.236
Walk Rate	0.364	0.366	-0.00183	0.00438	-0.417
Runs Allowed Per 9 Innings	4.900	5.037	-0.137*	0.0535	-2.559
Innings Pitched	132.382	136.038	-3.656	2.112	-1.731

Sample Size: Control Group = 2,344; Transition Group = 2,326
*** p<0.001, ** p<0.01, * p<0.05

Table 4.3: Independent Two-Sample t-test: Transition Group - Early Career Entry-Level Starting Pitcher to Relief Pitcher Transitions

Variable	Control Group Mean	Transition Group Mean	Difference	Standard Error	t-value
Hit Rate	1.010	1.050	-0.0400***	0.00649	-6.166
Strikeout Rate	0.799	0.756	0.0437***	0.00623	7.010
Walk Rate	0.456	0.439	0.0172***	0.00469	3.663
Runs Allowed Per 9 Innings	4.955	5.253	-0.297***	0.0559	-5.321
Innings Pitched	41.783	55.057	-13.274***	0.670	-19.822

Sample Size: Control Group = 7,770; Transition Group = 2,608
*** p<0.001, ** p<0.01, * p<0.05

Our t-tests yield results that contradict the theory that entry-level relief pitchers are

systematically lower ability pitchers. In both cases, the pitchers who transition roles perform significantly worse than their comparison group for hit rate, strikeout rate, and runs allowed per 9 innings. Pitchers who transition roles do not perform significantly worse than the control group for walk rate and innings pitched in both cases. These are striking results given that entry-level starting pitchers maintained survival advantages over entry-level relief pitchers for both comparison sets. Thus, these results suggest that the survival advantage for entry-level starting pitchers cannot be explained by systematic differences in quality between the two entry-level pitcher groups.

One limitation of our Kaplan-Meier approach thus far is that it does not allow for multi-variate analysis in determining the variables that influence pitcher survival. To address this issue we utilize a Cox-Proportional Hazard model to simultaneously account for the impact of multiple covariates on the career length of MLB pitchers. This approach will allow for us to control for pitcher performance measures and will serve as a robustness check for our previous results based on our Kaplan-Meier estimates. The structure of the Cox-Proportional Hazard model is as follows:

(4.2) $$\lambda(t|X) \ = \ \lambda_0(t)exp(X\beta) \ + \ \epsilon,$$

where the equation represents the hazard function at time t given the covariate vector X and β represents the vector of coefficients to be estimated. Covariates included in our model are defined in the table below:

Table 4.4: Covariates for Cox-Proportional Hazard Model

Covariate	Definition
RP	Dummy variable indicating entry-level relief pitcher
Trans	Dummy variable indicating role transition within first three seasons of career
Age	Entry-level age
ZInnings	Entry-level innings pitched standardized by season and entry-level pitcher group
ZHitRate	Entry-level hit rate standardized by season and entry-level pitcher group
ZKRate	Entry-level strikeout rate standardized by season and entry-level pitcher group
ZBBRate	Entry-level walk rate standardized by season and entry-level pitcher group
Time	Time Trend

We standardize performance metrics by pitcher type and season to account for the structural differences in the roles between starting pitchers and relief pitchers as previous mentioned. For the stability of our rate metrics, we only include entry-level seasons in which the pitcher had at least five innings pitched once again. Additionally, we interact our entry-level relief pitcher dummy variable (RP) with all of our other covariates to determine if significant differences exist in the impact these variables have on the career length of entry-level relief pitchers compared to entry-level starting pitchers.

In regards to our performance variables, if structural differences exist in the overall quality of starting pitchers and relief pitchers, we would expect to see significant coefficients for the interaction terms between our relief pitcher dummy variable and performance variables. In particular, if lower quality pitchers structurally get sorted into the entry-level reliever pool, we expect that relief pitchers will need to perform at a level higher on the standardized performance distribution than entry-level starting pitchers to achieve a similar career length. Therefore, a one standard deviation improvement for any of our performance metrics should be significantly less valuable in increasing career length for entry-level relief pitchers compared to entry-level starting pitchers if a systematic difference in quality exists between the two pitcher types.

Once again, as a robustness check we run our Cox-Proportional Hazard model using different thresholds for our starting pitcher vs. relief pitcher classification. The first model follows the primary classification we have used thus far and classifies a pitcher as a relief pitcher if he makes at least 50% of his appearances as a relief pitcher and as a starting pitcher otherwise. The second model narrows the starting pitcher classification while widening the relief pitcher classification by classifying a pitcher as a relief pitcher if he makes at least 25% of his appearances as a relief pitcher and as a starting pitcher otherwise. The third model widens the starting pitcher classification while narrowing the relief pitcher classification by classifying a pitcher as a relief pitcher if he makes at least 75% of his appearances as a relief pitcher and as a starting pitcher otherwise. In addition to the covariates listed in Table 4.4, we also include team level fixed effects in all models as well. The results of our Cox-Proportional Hazard models are included below:

Table 4.5: Cox-Proportional Hazard Model Coefficients

Variable	RP: 50%	RP: 25%	RP: 75%
RP	1.244***	1.158***	1.258***
	(0.414)	(0.436)	(0.398)
Trans	0.00609	-0.0568	-0.0143
	(0.0636)	(0.0731)	(0.0596)
ZHitRate	0.00622	-0.00856	0.0134
	(0.0257)	(0.0294)	(0.0230)
ZKRate	-0.127***	-0.140***	-0.136***
	(0.0292)	(0.0339)	(0.0254)
ZBBRate	0.0691***	0.0681**	0.0605**
	(0.0266)	(0.0297)	(0.0242)
ZInnings	-0.154***	-0.165***	-0.183***
	(0.0488)	(0.0530)	(0.0448)
Age	0.165***	0.168***	0.171***
	(0.0143)	(0.0161)	(0.0128)
Time	-0.00180	-0.000356	-0.00383
	(0.00292)	(0.00339)	(0.00258)
RP*Trans	-0.484***	-0.357***	-0.415***
	(0.0914)	(0.0953)	(0.0922)
RP*ZHitRate	0.0471	0.0588*	0.0441
	(0.0324)	(0.0348)	(0.0313)
RP*ZKRate	-0.000878	0.0204	0.0210
	(0.0363)	(0.0396)	(0.0344)
RP*ZBBRate	0.0153	0.0205	0.0278
	(0.0335)	(0.0354)	(0.0324)
RP*ZInnings	-0.0179	0.000324	0.00370
	(0.0581)	(0.0608)	(0.0557)
RP*Age	-0.0371**	-0.0326*	-0.0422***
	(0.0172)	(0.0184)	(0.0164)
RP*Time	0.00228	0.000990	0.00496
	(0.00349)	(0.00380)	(0.00334)
Observations	3,456	3,456	3,456
Failures	3,172	3,172	3,172
Log Likelihood	-22740.163	-22747.200	-22746.511

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Team level fixed effects included

Our estimates confirm several of our prior findings. The relief pitcher indicator variable is positive and significant for all models, indicating that being an entry-level relief pitcher has a negative overall impact on career length. We see that our standardized strikeout rate,

walk rate and innings pitched variables are all significant and with the expected signs in each model. Additionally, as expected, players with higher entry-level ages tend to have significantly shorter career lengths. As for our relief pitcher interaction terms, we observe that an early career role transition for entry-level relief pitchers significantly improves their career length while an early career role transition for entry-level starting pitchers has an insignificant impact on their career lengths. The only other variable consistently significant at the .05 level when interacted with our relief pitcher dummy variable is age. These results indicate that entering the league at a later age is more detrimental to the survival of entry-level starting pitchers compared to entry-level relief pitchers. We do not observe any other significant coefficients across models for our relief pitcher interaction terms, including all performance variables. This supports our previous finding that survival differences cannot be explained by systematic differences in quality between starting pitchers and relief pitchers.

Thus far we have focused exclusively on career length as our dependent variable of interest in our analysis showing that an entry-level reliever assignment has a negative impact on one's career compared to an entry-level starting pitcher assignment. As an alternate specification we utilize the model structure from our Cox-Proportional Hazard models to determine if our independent variables have a similar effect on career earnings. This specification will allow us to account for any compensation advantages or disadvantages for these roles that career length may not capture. For this specification we use an OLS model with team level fixed effects and standard errors robust to heteroskedasticity with the natural log of career earnings as our dependent variable and the same covariates from Table 4.4 as our independent variables. The salary data we use to measure career earnings is obtained from the Lahman Database. Since the salary data only includes guaranteed salaries for Major League contracts, players that only earned salaries on prorated contracts at the Major League level after being promoted from the minor leagues are excluded in this data. Once again three separate models are run using the exact same respective relief pitcher and starting pitcher classifications for

the results in Table 4.5. The results of our career earnings regressions are included below.

Table 4.6: Career Earnings Regressions - Dependent Variable: ln(Career Earnings)

Variable	RP: 50%	RP: 25%	RP: 75%
RP	-2.435***	-2.501**	-2.491***
	(0.927)	(1.017)	(0.899)
Trans	-0.425***	-0.514***	-0.331***
	(0.127)	(0.149)	(0.120)
ZHitRate	-0.0586	-0.0569	-0.0497
	(0.0545)	(0.0614)	(0.0505)
ZKRate	0.275***	0.252***	0.317***
	(0.0612)	(0.0742)	(0.0554)
ZBBRate	-0.193***	-0.144**	-0.178***
	(0.0550)	(0.0625)	(0.0507)
ZInnings	0.493***	0.517***	0.531***
	(0.0919)	(0.0991)	(0.0854)
Age	-0.189***	-0.198***	-0.204***
	(0.0337)	(0.0392)	(0.0312)
Time	0.0104*	0.00663	0.0153***
	(0.00629)	(0.00723)	(0.00567)
RP*Trans	1.225***	1.279***	1.076***
	(0.177)	(0.188)	(0.179)
RP*ZHitRate	-0.0355	-0.0227	-0.0469
	(0.0677)	(0.0728)	(0.0661)
RP*ZKRate	-0.0185	0.00393	-0.0956
	(0.0751)	(0.0848)	(0.0720)
RP*ZBBRate	0.0494	-0.0255	0.0382
	(0.0702)	(0.0750)	(0.0686)
RP*ZInnings	-0.147	-0.166	-0.152
	(0.107)	(0.113)	(0.102)
RP*Age	0.0456	0.0445	0.0614
	(0.0390)	(0.0433)	(0.0374)
RP*Time	0.00455	0.00557	-0.000834
	(0.00761)	(0.00825)	(0.00736)
Constant	20.50***	20.85***	20.54***
	(0.889)	(1.009)	(0.834)
Observations	2,171	2,171	2,171
R-squared	0.203	0.200	0.200

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Team level fixed effects included

These results are largely consistent with our Cox Proportional Hazard model results from Table 4.5. An entry-level relief pitcher assignment has a negative impact on career earnings

while standardized strikeout rate, walk rate and innings pitched variables are all significant and with the expected signs once again in each model. Older debut ages negatively impact career earnings while transitioning from a relief pitcher to starting pitcher early in one's career positively influence career earnings. One difference between the results in Table 4.5 and Table 4.6 is that entry-level starting pitchers who transition to relief pitchers early in their careers see a significant decrease in career earnings in all regressions, whereas the impact of this transition on career length was insignificant in Table 4.5. Once again, we do not observe any other significant coefficients across models for our relief pitcher interaction terms, including all performance variables. We take these results to be a confirmation of our previous findings on the impact of a pitcher's entry-level assignment on his career opportunities.

4.4   Discussion and Conclusions

This paper has conducted a survival analysis to determine if a pitcher's team-assigned, entry-level role impacts his ability to survive in the league. We find strong evidence of entry-level relief pitchers having a significant disadvantage in survival compared to entry-level starting pitchers. Our Kaplan-Meier estimates and log-rank test results show that the career lengths of entry-level relief pitchers are lower than that of entry-level starting pitchers and that these differences are statistically significant. Furthermore, after completing our analysis conditioning our sample on years of experience, we find that statistically significant survival differences between entry-level starting pitchers and relief pitchers last until a pitcher reaches 12 years of experience. Thus, the survival disadvantage for entry-level relief pitchers essentially lasts for the duration of a given pitcher's career, with the exception of entry-level relief pitchers who survive to the upper tail of the career length distribution.

This paper then uses the frequency of early-career role transitions to determine whether survival disadvantages (advantages) continue to exist for entry-level relief (starting) pitchers

given an early-career role transition. Our Kaplan-Meier estimates and log-rank test continue to show that entry-level relief pitchers who transitioned to starting pitchers early in their careers still face a statistically significant survival disadvantage compared to other entry-level starting pitchers, while entry-level starting pitchers who transitioned to relief pitchers early in their careers still maintain a statistically significant survival advantage compared to other entry-level relief pitchers. We also find that in both cases that pitchers who transition roles perform significantly worse than those who do not for hit rate, strikeout rate, and runs allowed per 9 innings. Thus, survival advantages for entry-level starting pitchers who transition roles early in their careers cannot be explained by differences in performance and systematic differences in pitcher quality between entry-level starting pitchers and entry-level relief pitchers.

Finally, to conduct a multivariate survival analysis we utilize a Cox-Proportional Hazard model using different thresholds for our pitcher type classification. The results confirm our prior findings of a significant survival disadvantage for entry-level relief pitchers while standardized strikeout rate, walk rate and innings pitched variables are all significant in determining a player's career length. Additionally, we do not find any significant differences in the impact of performance variables standardized by position on career length for starting pitchers versus relief pitchers. This further suggests that differences in survival between entry-level starting pitchers and entry-level relief pitchers cannot be explained by systematic differences in quality. As a robustness check we then used an OLS regression to study the impact of a pitcher's entry-level role assignment on his career earnings. The results are largely consistent with our observations from our Cox-Proportional Hazard model analysis as the same variables that tend to have a significant impact on career length tend to also have a significant impact on career earnings.

There are several potential reasons that explain why this survival difference between entry-level pitcher types exist. First, as illustrated in Figures 4.1 and 4.2, the growth of the

relief pitcher position in Major League Baseball has led to a much larger labor supply of players compared to any other position in MLB. As a result, we would expect to see greater competition for relief pitcher positions and lower tolerance for poor positions compared to other positions in baseball. Secondly, given that starting pitchers have historically been viewed as the more valuable pitcher type, teams may perceive a pitcher's entry-level position as a signal of a pitcher's ability, despite the fact that we do not find evidence of structural differences in ability between the two groups. Finally, differences in injury risk due to the structural differences in the roles of starting pitchers and relief pitchers may be another potential explanation of our findings. However, most of the previous research on pitcher injury suggests that injuries tend to have a greater impact on starting pitchers compared to relief pitchers since they tend to have a greater workload than relief pitchers in the form of innings pitched and pitches thrown (Gibson et al. [2007], Cerynik et al. [2008], Bradbury & Forman [2012], Jones et al. [2013], Karakolis et al. [2013], Whiteside et al. [2016], and Marshall et al. [2018]). Nevertheless, given that relief pitchers tend to pitch in far more games with fewer rest days between appearances compared to starting pitchers, an interesting area of future research would be to further quantify the difference in injury risk between these positions.

The results of this study also reveal how the early-career decisions made by teams on behalf of their players greatly influence the trajectory of player careers in Major League Baseball. Teams subjectively decide whether pitchers debut in the league as starting pitchers or relief pitchers, but our results show that this decision can dramatically affect a player's ability to survive in the league. Additionally, given that players often face restricted mobility early in their careers due to the reserve clause, players often do not have other employment options if they are unhappy with their team-assigned role unless they are granted a release by their current team. This essentially creates a situation for which pitchers are often forced to accept their entry-level roles, greatly influencing their survivability and career path in the

league.

While entry-level relief pitchers clearly face a survival disadvantage compared to entry-level starting pitchers, it should also be noted that there are surely certain situations for which an entry-level relief pitcher assignment may end up having a positive effect on career length. The most obvious example would be if a pitcher is given an opportunity to establish himself as a starting pitcher in the minor leagues, but struggles in the role due to poor performance. In such a case, the only realistic option for the pitcher to reach the Major Leagues may be as a relief pitcher, for which less stamina in a single game is required compared to starting pitchers. Thus, in such a situation, an entry-level relief pitcher role would be in the best interests of both the player and the team.

However, a potential issue arises when the best interests of the player and team differ. Such a situation would be when a pitcher's opportunity to establish himself in a starting pitcher role is either limited or absent (many pitchers are drafted out of college and assigned a relief pitcher role immediately in the minor leagues) prior to debuting at the Major League level. The problem also presents itself when the player is assigned an entry-level relief pitcher role due to the needs or roster construction of the team at the time of the pitcher's debut. Another interesting area of future research would be to investigate the determinants of a pitcher's entry-level role and transition likelihood at a more granular level. Since this paper does not find significant differences in pitcher quality between pitcher types, additional variables to consider may include more physiological attributes (height, weight, pitch velocity, etc.) or more structural attributes (pitcher role in college/high school, college/high school attended, team organizational depth at time of player draft, etc.).

This study illustrates how the early-career decisions made by employers on behalf of their employees can ultimately have career altering implications for workers. Given the results of this study, future research investigating how such decisions impact the career paths of their employees in other industries is warranted. Specifically in the sports industry, these results

raise issues with the current institutional rules of baseball pertaining to a team's control over a player once he is drafted by the team. Such findings suggest that changes to the collective bargaining agreement shrinking a team's control over minor league players may be warranted as well.

Chapter V

Conclusion

This dissertation utilizes the abundance of player, team, and wage data in Major League Baseball to conduct three studies pertaining to the player labor market and team decision making process in Major League Baseball.

The first study (Chapter 2) investigates whether the salary determination process for free agent players significantly changed after the publication of *Moneyball* while using run values to structurally determine a player's exact contributions to team production. The structural nature of run values improves upon the "ad-hoc" metrics used in previous research and allows us to more accurately assess a player's productivity as well as the impact of his performance outcomes on the salary determination process. It is found that the labor market for free agent players appropriately rewarded players according to their run value contributions both before and after publication of *Moneyball*, providing evidence contrary to the conclusions of previous research based on more "ad-hoc" measures of player productivity.

The second study (Chapter 3) uses a financial portfolio theory framework to investigate the impact of player and team mean and variance on team production. Using this framework, an analysis of trade efficiency is completed over the timeframe 1994-2016. It is found that both mean and variance positively affect team production. Given that variance is largely ignored in most measures of productivity in professional sports, this study illustrates the

need to account for performance variance when evaluating talent. This study also finds that trade efficiency has improved dramatically over this timeframe, consistent with the growth of analytics used in the decision-making process of professional sports teams during this time.

The final study (Chapter 4) investigates the impact of a pitcher's entry-level assignment on his ability to survive in the league. It is found that entry-level relief pitchers experience significantly shorter careers than entry-level starting pitchers and that significant survival differences continue to persist even when focusing exclusively on pitchers that transition roles early in their careers. Furthermore, this study does not find that entry-level relief pitchers are systematically of lower ability than entry-level starting pitchers, indicating that these differences in survival cannot be explained by differences in pitcher quality.

Each paper within this dissertation illustrates how the professional sports industry can be used to help contribute to the more general body of research in disciplines such as economics and finance. Thanks to the unique setting of professional sports, where productivity can be easily observed and quantified, as well as the abundance of publicly available performance and wage data, researchers are able to test theory and investigate research questions related to these broader disciplines. Such studies can be otherwise be difficult to conduct for other industries due to limitations in data availability and objective performance measurement. As a result, the areas of sport economics and sport finance should continue to provide fruitful research going forward. Areas of future study pertaining to these papers are self-contained within each respective chapter.

REFERENCES

Abid, F., Leung, P. L., Mroua, M., & Wong, W. K. (2014). International diversification versus domestic diversification: Mean-variance portfolio optimization and stochastic dominance approaches. Journal of Risk and Financial Management, 7(2), 45-66.

Adler, M., & Dumas, B. (1983). International portfolio choice and corporation finance: A synthesis. The Journal of Finance, 38(3), 925-984.

Ahlin, L., Andersson, M., & Thulin, P. (2014). Market thickness and the early labour market career of university graduates: An urban advantage?. Spatial Economic Analysis, 9(4), 396-419.

Albert, J. (2016) "Linear Weights on R" https://baseballwithr.wordpress.com/2016/06/13/linear-weights-on-r/

Albert, J., & Bennett, J. (2001). Curve ball: Baseball, statistics, and the role of chance in the game. Springer Science & Business Media.

Altonji, J. G. (2005). Employer learning, statistical discrimination and occupational attainment. American Economic Review, 95(2), 112-117.

Anderson, A. (2013). Trading and under-diversification. Review of Finance, 17(5), 1699-1741.

Arouri, M. E. H., & Nguyen, D. K. (2010). Oil prices, stock markets and portfolio investment: Evidence from sector analysis in Europe over the last decade. Energy policy, 38(8), 4528-4539.

Baker, J., Koz, D., Kungl, A. M., Fraser-Thomas, J., & Schorer, J. (2013). Staying at the top: playing position and performance affect career length in professional sport. High Ability Studies, 24(1), 63-76.

Barros, C. P., Frick, B., & Passos, J. (2009). Coaching for survival: The hazards of head coach careers in the German 'Bundesliga'. Applied Economics, 41(25), 3303-3311.

Basak, S., & Chabakauri, G. (2010). Dynamic mean-variance asset allocation. The Review of Financial Studies, 23(8), 2970-3016.

Baumer, B. S., Jensen, S. T., & Matthews, G. J. (2015). openWAR: An open source system for evaluating overall player performance in major league baseball. Journal of Quantitative Analysis in Sports, 11(2), 69-84.

Baumer, B., & Zimbalist, A. (2014). Quantifying market inefficiencies in the baseball players' market. Eastern Economic Journal, 40(4), 488-498.

Baumer, B., & Zimbalist, A. (2014). The sabermetric revolution: Assessing the growth of analytics in baseball. University of Pennsylvania Press.

Berri, D. J., & Bradbury, J. C. (2010). Working in the land of the metricians. Journal of Sports Economics, 11(1), 29-47.

Best, M. J., & Grauer, R. R. (1991). On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. The review of financial studies, 4(2), 315-342.

Best, M. J., & Grauer, R. R. (1991). Sensitivity analysis for mean-variance portfolio problems. Management Science, 37(8), 980-989.

Biemann, T., & Braakmann, N. (2013). The impact of international experience on objective and subjective career success in early careers. The International Journal of Human Resource Management, 24(18), 3438-3456.

Blass, A. A. (1992). Does the baseball labor market contradict the human capital model of investment?. The Review of Economics and Statistics, 261-268.

Bradbury, J. C. (2010). Hot stove economics: understanding baseball's second season. Springer Science & Business Media.

Bradbury, J. C. (2013). What is right with Scully estimates of a player's marginal revenue product. Journal of Sports Economics, 14(1), 87-96.

Bradbury, J. C., & Forman, S. L. (2012). The impact of pitch counts and days of rest on performance among major-league baseball pitchers. The Journal of Strength & Conditioning Research, 26(5), 1181-1187.

Brauneis, A., & Mestel, R. (2019). Cryptocurrency-portfolios in a mean-variance framework. Finance Research Letters, 28, 259-264.

Brewer, H. L. (1981). Investor benefits from corporate international diversification. Journal of Financial and Quantitative Analysis, 113-126.

Brown, D. T., Link, C. R., & Rubin, S. L. (2017). *Moneyball* After 10 Years: How Have Major League Baseball Salaries Adjusted? Journal of Sports Economics, 18(8), 771-786.

Butler, M. R. (1995). Competitive balance in major league baseball. The American Economist, 39(2), 46-52.

Campa, J. M., & Fernandes, N. (2006). Sources of gains from international portfolio diversification. Journal of Empirical Finance, 13(4-5), 417-443.

Cardozo, R. N., & Smith, D. K. (1983). Applying financial portfolio theory to product portfolio decisions: An empirical study. Journal of Marketing, 47(2), 110-119.

Cerynik, D. L., Ewald, T. J., Sastry, A., Amin, N. H., Liao, J. G., & Tom, J. A. (2008). Outcomes of isolated glenoid labral injuries in professional baseball pitchers. Clinical Journal of Sport Medicine, 18(3), 255-258.

Chatterjee, S., & Yilmaz, M. (1991). Parity in baseball: Stability of evolving systems?. Chance, 4(3), 37-42.

Chiu, C. H., & Choi, T. M. (2016). Supply chain risk analysis with mean-variance models: A technical review. Annals of Operations Research, 240(2), 489-507.

Chiu, M. C., & Wong, H. Y. (2011). Mean–variance portfolio selection of cointegrated assets. Journal of Economic Dynamics and Control, 35(8), 1369-1385.

Choi, T. M., Li, D., & Yan, H. (2008). Mean–variance analysis for the newsvendor problem. IEEE transactions on systems, man, and cybernetics-Part a: systems and humans, 38(5), 1169-1180.

Clarke, R., De Silva, H., & Thorley, S. (2011). Minimum-variance portfolio composition. The Journal of Portfolio Management, 37(2), 31-45.

Cochran-Smith, M., McQuillan, P., Mitchell, K., Terrell, D. G., Barnatt, J., D'Souza, L., ... & Gleeson, A. M. (2012). A longitudinal study of teaching practice and early career decisions: A cautionary tale. American educational research journal, 49(5), 844-880.

Coeurdacier, N., & Guibaud, S. (2011). International portfolio diversification is better than you think. Journal of international money and finance, 30(2), 289-308.

Cox, T. H., & Harquail, C. V. (1991). Career paths and career success in the early career stages of male and female MBAs. Journal of Vocational Behavior, 39(1), 54-75.

Das, S. R., & Uppal, R. (2004). Systemic risk and international portfolio choice. The Journal of Finance, 59(6), 2809-2834.

DeAngelis, K. J., Wall, A. F., & Che, J. (2013). The impact of preservice preparation and early career support on novice teachers' career intentions and decisions. Journal of teacher education, 64(4), 338-355.

DeBrock, L., Hendricks, W., & Koenker, R. (2004). Pay and performance: The impact of salary distribution on firm-level outcomes in baseball. Journal of Sports Economics, 5(3), 243-261.

Deli, D. (2013). Assessing the relative importance of inputs to a production function: getting on base versus hitting for power. Journal of Sports Economics, 14(2), 203-217.

Dossinger, K., Wanberg, C. R., Choi, Y., & Leslie, L. M. (2019). The beauty premium: The role of organizational sponsorship in the relationship between physical attractiveness and early career salaries. Journal of Vocational Behavior, 112, 109-121.

Dougherty, T. W., Dreher, G. F., & Whitely, W. (1993). The MBA as careerist: An analysis of early-career job change. Journal of management, 19(3), 535-548.

Driessen, J., & Laeven, L. (2007). International portfolio diversification benefits: Cross-country evidence from a local perspective. Journal of Banking & Finance, 31(6), 1693-1712.

Driessen, J., & Maenhout, P. (2007). An empirical portfolio perspective on option pricing anomalies. Review of Finance, 11(4), 561-603.

Ducking, J., & Bollinger, C. R. (2014). The effects of minimum salaries on career length: evidence from the National Football League (NFL). Available at SSRN 2466056.

Duquette, C. M., Cebula, R. J., & Mixon Jr, F. G. (2019). Major league baseball's Moneyball at age 15: a re-appraisal. Applied Economics, 51(52), 5694-5700.

Eckard, E. W. (2001). Baseball's blue ribbon economic report: Solutions in search of a problem. Journal of Sports Economics, 2(3), 213-227.

Einolf, K. W. (2004). Is winning everything? A data envelopment analysis of Major League Baseball and the National Football League. Journal of Sports Economics, 5(2), 127-151.

Eun, C. S., & Resnick, B. G. (1994). International diversification of investment portfolios: US and Japanese perspectives. Management science, 40(1), 140-161.

Fletcher, J. (2009). Risk Reduction and Mean-Variance Analysis: An Empirical Investigation. Journal of Business Finance & Accounting, 36(7-8), 951-971.

Flinkman, M., Isopahkala-Bouret, U., & Salanterä, S. (2013). Young registered nurses' intention to leave the profession and professional turnover in early career: a qualitative case study. International Scholarly Research Notices, 2013.

Fort, R. (1992). Pay and Performance: Is the Field of Dreams Barren?. Diamonds are forever: The business of baseball, 134-162.

Friedman, H. C. (1971). Real estate investment and portfolio theory. Journal of Financial and Quantitative Analysis, 861-874.

Frost, P. A., & Savarino, J. E. (1986). An empirical Bayes approach to efficient portfolio selection. Journal of Financial and Quantitative Analysis, 293-305.

Gârleanu, N., & Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. The Journal of Finance, 68(6), 2309-2340.

Gebel, M. (2010). Early career consequences of temporary employment in Germany and the UK. Work, employment and society, 24(4), 641-660.

Gibson, B. W., Webner, D., Huffman, G. R., & Sennett, B. J. (2007). Ulnar collateral ligament reconstruction in major league baseball pitchers. The American journal of sports medicine, 35(4), 575-581.

Gilmore, C. G., & McManus, G. M. (2002). International portfolio diversification: US and Central European equity markets. Emerging Markets Review, 3(1), 69-83.

Goldacre, M. J., Laxton, L., & Lambert, T. W. (2010). Medical graduates' early career choices of specialty and their eventual specialty destinations: UK prospective cohort studies. Bmj, 341.

Gould, S. J. (1986). Entropic homogeneity isn't why no one hits. 400 any more. Discover, 7(8), 60-66.

Groothuis, P. A., & Hill, J. R. (2004). Exit discrimination in the NBA: A duration analysis of career length. Economic Inquiry, 42(2), 341-349.

Györfi, L., Udina, F., & Walk, H. (2008). Nonparametric nearest neighbor based empirical portfolio selection strategies. Statistics & Risk Modeling, 26(2), 145-157.

Hakes, J. K., & Sauer, R. D. (2006). An economic evaluation of the *Moneyball* hypothesis. Journal of Economic Perspectives, 20(3), 173-186.

Hakes, J. K., & Sauer, R. D. (2007). The *Moneyball* anomaly and payroll efficiency: A further investigation. International Journal of Sport Finance, 2(4), 177.

Hakes, J. K., & Turner, C. (2011). Pay, productivity and aging in Major League Baseball. Journal of Productivity Analysis, 35(1), 61-74

Hall, S., Szymanski, S., & Zimbalist, A. S. (2002). Testing causality between team performance and payroll: the cases of Major League Baseball and English soccer. Journal of Sports Economics, 3(2), 149-168.

Hickman, K. A., Teets, W. R., & Kohls, J. J. (1999). Social investing and modern portfolio theory. American Business Review, 17(1), 72.

Hjalmarsson, E., & Manchev, P. (2012). Characteristic-based mean-variance portfolio choice. Journal of Banking & Finance, 36(5), 1392-1401.

Hoang, H., & Rascher, D. (1999). The NBA, exit discrimination, and career earnings. Industrial Relations: A Journal of Economy and Society, 38(1), 69-91.

Holmes, P. (2011). New evidence of salary discrimination in major league baseball. Labour Economics, 18(3), 320-331.

Holmes, P. M., Simmons, R., & Berri, D. J. (2018). *Moneyball* and the Baseball Players' Labor Market. International Journal of Sport Finance, 13(2).

Humphreys, B. R., & Pyun, H. (2017). Monopsony exploitation in professional sport: Evidence from Major League Baseball position players, 2000–2011. Managerial and Decision Economics, 38(5), 676-688.

Hussein, K. (2004). Ethical investment: empirical evidence from FTSE Islamic index. Islamic Economic Studies, 12(1).

Jensen, M. C., Black, F., & Scholes, M. S. (1972). The capital asset pricing model: Some empirical tests.

Jones, K. J., Conte, S., Patterson, N., ElAttrache, N. S., & Dines, J. S. (2013). Functional outcomes following revision ulnar collateral ligament reconstruction in Major League Baseball pitchers. Journal of shoulder and elbow surgery, 22(5), 642-646.

Jorion, P. (1985). International portfolio diversification with estimation risk. Journal of Business, 259-278.

Kahn, L. M. (1993). Free agency, long-term contracts and compensation in Major League Baseball: Estimates from panel data. The Review of Economics and Statistics, 157-164.

Kahn, L. M. (2000). The sports business as a labor market laboratory. Journal of economic perspectives, 14(3), 75-94.

Karakolis, T., Bhan, S., & Crotin, R. L. (2013). An inferential and descriptive statistical examination of the relationship between cumulative work metrics and injury in Major League Baseball pitchers. The Journal of Strength & Conditioning Research, 27(8), 2113-2118.

Kelley, J. (1978). Wealth and family background in the occupational career: Theory and cross-cultural data. British Journal of Sociology, 94-109.

Koji, Y., & Nakazono, Y. (2020). Moneyball Revisited: Some Counter-Evidence. Available at SSRN 3757785.

Kourtidis, D., Šević, Ž., & Chatzoglou, P. (2011). Investors' trading activity: A behavioural perspective and empirical results. The Journal of Socio-Economics, 40(5), 548-557.

Krautmann, A. C. (1999). What's wrong with Scully-estimates of a player's marginal revenue product. Economic inquiry, 37, 369-381.

Krautmann, A. C. (2017). Risk-averse team owners and players' salaries in Major League Baseball. Journal of Sports Economics, 18(1), 19-33.

Krautmann, A. C., & Donley, T. D. (2009). Shirking in major league baseball revisited. Journal of Sports Economics, 10(3), 292-304.

Krautmann, A. C., Gustafson, E., & Hadley, L. (2003). A note on the structural stability of salary equations: major league baseball pitchers. Journal of Sports Economics, 4(1), 56-63.

Krautmann, A. C., & Oppenheimer, M. (2002). Contract length and the return to performance in Major League Baseball. Journal of Sports Economics, 3(1), 6-17.

Krautmann, A. C., Von Allmen, P., & Berri, D. (2009). The underpayment of restricted players in North American sports leagues. International Journal of Sport Finance, 4(3), 161.

Kroll, Y., Levy, H., & Markowitz, H. M. (1984). Mean-variance versus direct utility maximization. The Journal of Finance, 39(1), 47-61.

Lai, T. L., Xing, H., & Chen, Z. (2011). Mean–variance portfolio optimization when means and covariances are unknown. The Annals of Applied Statistics, 5(2A), 798-823.

Lee, C. K., & Wingreen, S. C. (2010). Transferability of knowledge, skills, and abilities along IT career paths: An agency theory perspective. Journal of Organizational Computing and Electronic Commerce, 20(1), 23-44.

Lessard, D. R. (1973). International portfolio diversification: a multivariate analysis for a group of Latin American countries. The Journal of Finance, 28(3), 619-633.

Levy, H., & Lerman, Z. (1988). The benefits of international diversification in bonds. Financial Analysts Journal, 44(5), 56-64.

Levy, H., & Sarnat, M. (1970). International diversification of investment portfolios. The American Economic Review, 60(4), 668-675.

Li, D., & Ng, W. L. (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. Mathematical finance, 10(3), 387-406.

Lindsey, G. R. (1963). An investigation of strategies in baseball. Operations Research, 11(4), 477-501.

Link, C. R., & Yosifov, M. (2012). Contract length and salaries compensating wage differentials in Major League Baseball. Journal of Sports Economics, 13(1), 3-19.

Lindsey, G. R. (1963). An investigation of strategies in baseball. Operations Research, 11(4), 477-501.

Lo, A. W., & Wang, J. (2000). Trading volume: definitions, data analysis, and implications of portfolio theory. The Review of Financial Studies, 13(2), 257-300.

Lubatkin, M., & Chatterjee, S. (1994). Extending modern portfolio theory into the domain of corporate diversification: does it apply?. Academy of Management Journal, 37(1), 109-136.

Maccheroni, F., Marinacci, M., Rustichini, A., & Taboga, M. (2009). Portfolio selection with monotone mean-variance preferences. Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics, 19(3), 487-521.

MacDonald, D. N., & Reynolds, M. O. (1994). Are baseball players paid their marginal products?. Managerial and Decision Economics, 15(5), 443-457.

MacKinlay, A. C., & Richardson, M. P. (1991). Using generalized method of moments to test mean-variance efficiency. The Journal of Finance, 46(2), 511-527.

Marburger, D. R. (1994). Bargaining power and the structure of salaries in major league baseball. Managerial and Decision Economics, 15(5), 433-441.

Marchi, M., & Albert, J. (2013). Analyzing baseball data with R. CRC Press.

Markowitz, H.M. (1952) Portfolio Selection. Journal of Finance, 7, 77-91.

Markowitz, H.M. (1959) Portfolio Selection: Efficient Diversification of Investments. John Wiley & Sons, New York.

Marshall, N. E., Jildeh, T. R., Okoroha, K. R., Patel, A., Moutzouros, V., & Makhni, E. C. (2018). Implications of core and hip injuries on Major League Baseball pitchers on the disabled list. Arthroscopy: The Journal of Arthroscopic & Related Surgery, 34(2), 473-478.

Mattox II, J. R., & Jinkerson, D. L. (2005). Using survival analysis to demonstrate the effects of training on employee retention. Evaluation and Program Planning, 28(4), 423-430.

Merton, R. C., Scholes, M. S., & Gladstein, M. L. (1978). The returns and risk of alternative call option portfolio investment strategies. Journal of Business, 183-242.

Miller, P. A. (2000). A theoretical and empirical comparison of free agent and arbitration-eligible salaries negotiated in major league baseball. Southern Economic Journal, 87-104.

Monti, H., Stinson, M., & Zehr, L. (2020). How Long Do Early Career Decisions Follow Women? The Impact of Employer History on the Gender Wage Gap. Journal of Labor Research, 41(3).

Musset, P., & Kurekova, L. M. (2018). Working it out: Career guidance and employer engagement.

Nagengast, A. J., Braun, D. A., & Wolpert, D. M. (2011). Risk-sensitivity and the mean-variance trade-off: decision making in sensorimotor control. Proceedings of the Royal Society B: Biological Sciences, 278(1716), 2325-2332.

Palczewski, A., & Palczewski, J. (2014). Theoretical and empirical estimates of mean–variance portfolio sensitivity. European Journal of Operational Research, 234(2), 402-410.

Petersen, A. M., Jung, W. S., Yang, J. S., & Stanley, H. E. (2011). Quantitative and empirical demonstration of the Matthew effect in a study of career longevity. Proceedings of the National Academy of Sciences, 108(1), 18-23.

Phylaktis, K., & Ravazzolo, F. (2005). Stock market linkages in emerging markets: implications for international portfolio diversification. Journal of International Financial Markets, Institutions and Money, 15(2), 91-106.

Platanakis, E., & Urquhart, A. (2020). Should investors include bitcoin in their portfolios? A portfolio theory approach. The British Accounting Review, 52(4), 100837.

Pliska, S. R. (1986). A stochastic calculus model of continuous trading: optimal portfolios. Mathematics of operations research, 11(2), 371-382.

Porter, R. B. (1973). An empirical comparison of stochastic dominance and mean-variance portfolio choice criteria. Journal of Financial and Quantitative Analysis, 587-608.

Quartz, K. H., Thomas, A., Anderson, L., Masyn, K., Lyons, K. B., & Olsen, B. (2008). Careers in motion: A longitudinal retention study of role changing among early-career urban educators. Teachers College Record, 110(1), 218-250.

Raaum, O., & Røed, K. (2006). Do business cycle conditions at the time of labor market entry affect future employment prospects?. The review of economics and statistics, 88(2), 193-210.

Raimondo, H. J. (1983). Free agents' impact on the labor market for baseball players. Journal of Labor Research, 4(2), 183-193.

Richards, D. G., & Guell, R. C. (1998). Baseball success and the structure of salaries. Applied Economics Letters, 5(5), 291-296.

Richardson, H. R. (1989). A minimum variance result in continuous trading portfolio optimization. Management Science, 35(9), 1045-1055.

Roques, F. A., Newbery, D. M., & Nuttall, W. J. (2008). Fuel mix diversification incentives in liberalized electricity markets: A Mean–Variance Portfolio theory approach. Energy Economics, 30(4), 1831-1849.

Ruefli, T. W. (1990). Mean-variance approaches to risk-return relationships in strategy: Paradox lost. Management Science, 36(3), 368-380.

Rugman, A. M. (1976). Risk reduction by international diversification. Journal of International Business Studies, 7(2), 75-80.

Salaga, S., Mills, B. M., & Tainsky, S. (2020). Employer-Assigned Workload and Human Capital Deterioration: Evidence From the National Football League. Journal of Sports Economics, 21(6), 628-659.

Schmidt, M. B. (2020). Labor demographics and productivity: all-star roster turnover and foreigners. Journal of Economic Studies.

Schmidt, M. B., & Berri, D. J. (2003). On the evolution of competitive balance: The impact of an increasing global search. Economic Inquiry, 41(4), 692-704.

Schmidt, M. B., & Berri, D. J. (2005). Concentration of playing talent: evolution in Major League Baseball. Journal of Sports Economics, 6(4), 412-419.

Schmillen, A., & Umkehrer, M. (2017). The scars of youth: Effects of early-career unemployment on future unemployment experience. International Labour Review, 156(3-4), 465-494.

Scully, G. W. (1974). Pay and performance in major league baseball. The American Economic Review, 64(6), 915-930.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. The journal of finance, 19(3), 425-442.

Sharpe, W. F. (1966). Mutual fund performance. The Journal of business, 39(1), 119-138.

Shawky, H. A., Kuenzel, R., & Mikhail, A. D. (1997). International portfolio diversification: a synthesis and an update. Journal of international financial markets, Institutions and Money, 7(4), 303-327.

Sommers, P. M., & Quinton, N. (1982). Pay and performance in major league baseball: The case of the first family of free agents. The Journal of Human Resources, 17(3), 426-436.

Stulz, R. M. (1995). International portfolio choice and asset pricing: An integrative survey. Handbooks in Operations Research and Management Science, 9, 201-223.

Stumpf, S. A., Doh, J. P., & Tymon Jr, W. G. (2010). The strength of HR practices in India and their effects on employee career success, performance, and potential. Human resource management: Published in Cooperation with the school of business administration, the university of Michigan and in alliance with the society of human resources management, 49(3), 353-375.

Syriopoulos, T. (2004). International portfolio diversification to Central European stock markets. Applied Financial Economics, 14(17), 1253-1268.

Tango, T. M., Lichtman, M. G., & Dolphin, A. E. (2006). The book: Playing the percentages in baseball. Potomac Books, Inc..

Teller, J., & Kock, A. (2013). An empirical investigation on how portfolio risk management influences project portfolio success. International Journal of Project Management, 31(6), 817-829.

Thorn, J., Palmer, P., & Reuther, D. (1985). The hidden game of baseball: A revolutionary approach to baseball and its statistics (Rev. Ed.).

Tobin, J. (1958). Liquidity preference as behavior towards risk. The review of economic studies, 25(2), 65-86.

Turocy, T. L. (2005). Offensive performance, omitted variables, and the value of speed in baseball. Economics Letters, 89(3), 283-286.

Van der Heijden, B. I., de Lange, A. H., Demerouti, E., & Van der Heijde, C. M. (2009). Age effects on the employability–career success relationship. Journal of vocational behavior, 74(2), 156-164.

Von Wachter, T., & Bender, S. (2006). In the right place at the wrong time: The role of firms and luck in young workers' careers. American Economic Review, 96(5), 1679-1705.

Whiteside, D., Martini, D. N., Lepley, A. S., Zernicke, R. F., & Goulet, G. C. (2016). Predictors of ulnar collateral ligament reconstruction in Major League Baseball pitchers. The American journal of sports medicine, 44(9), 2202-2209.

Williamson, I. O., & Cable, D. M. (2003). Predicting early career research productivity: The case of management faculty. Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior, 24(1), 25-44.

Winston, W. L. (2012). Mathletics: How gamblers, managers, and sports enthusiasts use mathematics in baseball, basketball, and football. Princeton University Press.

Witnauer, W. D., Rogers, R. G., & Saint Onge, J. M. (2007). Major league baseball career length in the 20th century. Population research and policy review, 26(4), 371-386.

Woolner, K. (2007). "Introduction to VORP: Value Over Replacement Player". https://web.archive.org/web/20070928064958/http://www.stathead.com/bbeng/woolner/vorpdescnew.htm

Wu, J., Li, J., Wang, S., & Cheng, T. C. (2009). Mean–variance analysis of the newsvendor model with stockout cost. Omega, 37(3), 724-730.

Wu, G., Wei, Y. M., Fan, Y., & Liu, L. C. (2007). An empirical analysis of the risk of crude oil imports in China using improved portfolio approach. Energy policy, 35(8), 4190-4199.

Yorke, D. A., & Droussiotis, G. (1994). The use of customer portfolio theory: An empirical survey. Journal of Business & Industrial Marketing.

Yu, J., & Yuan, Y. (2011). Investor sentiment and the mean–variance relation. Journal of Financial Economics, 100(2), 367-381.

Zech, C. E. (1981). An empirical estimation of a production function: The case of Major League Baseball. The American Economist, 25(2), 19-23.

Zhang, S., Zhao, T., & Xie, B. C. (2018). What is the optimal power generation mix of China? An empirical analysis using portfolio theory. Applied Energy, 229, 522-536.

Zhang, Y. J., & Chen, M. Y. (2018). Evaluating the dynamic performance of energy portfolios: Empirical evidence from the DEA directional distance function. European Journal of Operational Research, 269(1), 64-78.

Zhao, L., & Palomar, D. P. (2018). A markowitz portfolio approach to options trading. IEEE Transactions on Signal Processing, 66(16), 4223-4238.

Zimbalist, A. S. (1992). Baseball and billions: A probing look inside the big business of our national pastime. Basic Books.

Zimbalist, A. (2010). May the best team win: Baseball economics and public policy. Brookings Institution Press.