Process flexibility in baseball: The value of positional flexibility

Timothy C. Y. Chan, Douglas Fearing

Version post-print

Citation (published version) Chan, T. C. Y., & Fearing, D. (2019). Process flexibility in baseball: The value of positional flexibility. Management Science, 65(4), 1642–1666. 10.1287/mnsc.2017.3004

Submitted to *Management Science*
manuscript

Process flexibility in baseball: The value of positional flexibility

Timothy C. Y. Chan
Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, ON, M5S 3G8, Canada, tcychan@mie.utoronto.ca

Douglas Fearing
Research and Development, Los Angeles Dodgers, 1000 Vin Scully Avenue, Los Angeles, CA, 90012

This paper introduces the formal study of process flexibility to the novel domain of sports analytics. In baseball, *positional flexibility* is the analogous concept to process flexibility from manufacturing. We study the flexibility of players (plants) on a baseball team who produce innings-played at different positions (products). We develop models and metrics to evaluate expected and worst-case performance under injury risk (capacity uncertainty) with continuous player-position capabilities. Using Major League Baseball data, we quantify the impact of flexibility on team and individual performance, and explore the player chains that arise when injuries occur. We discover that top teams can attribute at least one to two wins per season to flexibility alone, generally due to long subchains in the infield or outfield. The least robust teams to worst-case injury, those whose performance is driven by one or two star players, are over four times as fragile as the most robust teams. We evaluate several aspects of individual flexibility, such as how much value individual players bring to their team in terms of average and worst-case performance. Finally, we demonstrate the generalizability of our framework for player evaluation by quantifying the value of potential free agent additions and uncovering the true "MVP" of a team.

Key words: baseball, analytics, positional flexibility, process flexibility, capacity uncertainty

1. Introduction

In this paper, we examine a long-standing, open question in baseball analytics: what is the value of positional flexibility? To illustrate the concept of positional flexibility, which refers to the ability of a player to play multiple positions, we highlight a Major League Baseball (MLB) game between the Tampa Bay Rays and the Seattle Mariners on May 2, 2012. Prior to the start of the sixth inning, B. J. Upton, the Rays' starting center fielder, experienced tightness in his right quadriceps.

As a precautionary measure, the Rays decided to remove Upton from the game. The following is taken from a news story posted on MLB.com after the game.[1]

> *Upton's departure set into motion a radical change of positions in the field for the Rays. Desmond Jennings shifted from left field to center, Matt Joyce went from right to left and Ben Zobrist moved from second base to right field. Elliot Johnson entered the game to play shortstop and hit in Upton's slot in the lineup. Sean Rodriguez moved from shortstop to third base and Will Rhymes shifted from third base to second.*

After this series of swaps, only the pitcher, catcher, and first baseman were left at their original positions. (See the Appendix for a brief introduction to the game of baseball.)

The concept of positional flexibility is well-known in baseball, though not necessarily well-understood especially when it comes to quantifying its value. In fact, the challenge of quantifying the value of positional flexibility was deemed one of baseball's 23 "Hilbert problems" (Woolner 2000), in reference to mathematician David Hilbert's 1900 address to the International Congress of Mathematicians and subsequent publication outlining 23 major unsolved problems in mathematics. Baseball's Hilbert problems include topics in defense, offense, pitching, player development, economics, strategic decisions, and tactical decisions. With respect to positional flexibility, the author writes:

> *A player who plays two positions at a league-average level gives his manager flexibility, both in setting up the team's roster and using in-game strategies. . . . Because roster spots are scarce, a team gets value from a player's ability to play multiple positions, but we do not yet have an understanding of how much value there is to having [such a player] on your roster.*

We believe this paper takes an important step forward in addressing this open question.

Positional flexibility is exhibited by non-pitching players (i.e., *position players*). According to our data, between 2007 and 2011 approximately 18% of the innings assigned to MLB position players were at non-primary positions, where the primary position is defined as the position the player spent the plurality of his time. Over the course of a season, such flexibility plays a critical role in setting the team's starting lineup. For example, major league position players missed a total of 11,184 days due to injury in 2013, which translates to roughly 373 days per team (Zimmerman 2013). The maximum was 1,146 days (New York Yankees) and the minimum was 109 days (Colorado Rockies). Injury risk is a major source of uncertainty and flexibility adds value by extending the set of options available to a team in the case of injury. This issue will be the focus of our paper, namely the role and value of flexibility in mitigating the negative effects of injuries.

Positional flexibility is the baseball analogue of manufacturing process flexibility. In particular, a baseball team can be viewed as a production network in which players (plants) produce

[1] http://mlb.mlb.com/news/article.jsp?ymd=20120502&content_id=30243696&vkey=news_tb&c_id=tb

innings-played to satisfy the demand for all positions (products) on a team. To help establish the connection for readers who may be unfamiliar with baseball but familiar with process flexibility in manufacturing, Table 1 summarizes what we believe to be an intuitive correspondence between the relevant concepts in these two domains.

Table 1 Conceptual correspondence between manufacturing and baseball.

Manufacturing	Baseball
Plants	Position players on the roster
Products	Non-pitching positions
Capability	Player skill (offense + defense)
Demand for each product	Total innings for each position
Plant capacity uncertainty	Player injury risk
Excess capacity in network	Roster depth

Although positional flexibility provides a team with options, flexibility alone is not enough. Teams also need *roster depth* and *capability*. Roster depth in the form of bench players (i.e., excess capacity) is needed when injuries occur. Capability refers to the ability to play positions with skill. Note that flexibility, roster depth, and capability are intertwined. Positional flexibility is useless without roster depth, and flexible players who play secondary positions poorly may offer little value. Furthermore, the interplay between a player's offensive and defensive contributions adds complexity to the evaluation of flexibility. For example, a player who plays multiple defensive positions well but is a poor hitter may actually be a worse option than simply having a superstar hitter with limited defensive capabilities.

Our methodological approach considers two extensions to the standard framework for studying manufacturing flexibility. First, previous work has mostly focused on "on-off" flexibility, where links in the network are either present or not. In this paper, we use the term "capability" as the continuous generalization of the on-off flexibility idea. For example, a player should be highly capable at his primary position, but generally will be less capable to varying degrees at other positions. Second, in addition to examining flexibility from an expected value context, we also introduce a worst-case (i.e., robust optimization) model and accompanying metrics. For example, a team may care more about optimizing their worst-case chances of making the playoffs, rather than maximizing their expected performance during the regular season. From the injury perspective, teams need insight into both their expected performance and worst-case performance. Teams that are susceptible to injuries to one or two key players on their roster may do well to improve their overall performance by identifying and training suitable back-up players, or acquiring back-up capacity through trades or the free agent market.

Leveraging the optimization frameworks and metrics we develop, we will pose and answer several questions that shed new light on the topic of positional flexibility in baseball. Our findings will be useful for managers, players, and fans of the sport. The first set of questions aims to uncover the fundamental value of positional flexibility at a team level. They aim to quantify effects that have not been quantified rigorously before. In particular, the questions are:

1. **What is the average value of flexibility at the team level? (Section 4.1)** Is it a material amount, e.g., the difference between making and not making the playoffs? Additionally, what are the characteristics of teams that derive much/little value from flexibility?

2. **How do we measure the robustness of a team to worst-case injury? (Section 4.2)** What are the characteristics of robust/fragile teams in terms of the players they possess?

The focus of the second set of analyses is on individual players. As such, these questions demonstrate how our methodological framework can provide managers with prescriptions for action regarding roster construction, and deeper insight into player performance and valuation, which may be useful for both managers and players in contract negotiations, for example. Such insight is likely of interest to fans as well, who could use our results to form new types of player rankings.

3. **What is the average value of flexibility at the individual level? (Section 4.3)** Can our framework identify the "Most Flexible Player", i.e., the player who adds the most value to his team from flexibility? What impact does a player's flexibility have on his own performance?

4. **Is it possible to improve team robustness to worst-case injury using only the existing roster? (Section 4.4)** How can we identify and reinforce the "weak links"?

5. **Given a list of free agents, which one should a team acquire to improve its performance the most? (Section 4.5)** How consistent is a player's added value across teams? Is there a trade-off between average and worst-case performance improvement?

6. **Who is the Most Valuable Player? (Section 4.6)** Can our flexibility framework be extended to measure the overall value of a player to his team, not just his flexibility value? Is a player's observable production representative of his true value to his team?

Our specific contributions in this paper are as follows:

- We introduce the formal study of flexibility in the novel problem domain of baseball. We evaluate the flexibility of Major League Baseball teams in both expected value and worst-case contexts, the contribution of flexible players to team performance, and demonstrate how our approach can help managers identify and address vulnerabilities in their team's composition. In the process, we address a long-standing, open question in baseball analytics.

- We develop a novel robust (min-max) optimization model to analyze the impact of flexibility in a worst-case setting. We present a bilinear formulation of our min-max model and derive an equivalent mixed-integer programming formulation that solves efficiently. We emphasize that our goal is to quantify the value of flexibility at a structural level (min-max) rather than to develop a roster that is robust to injury (max-min).

- We present several simple yet general metrics to quantify the value of flexibility in expected value and worst-case settings using our continuous capability representation of flexibility. In our computational experiments we apply our models and metrics to derive baseball-specific insights. Specifically, we find the following answers to the questions posed above:

1. Top teams can attribute at least one to two wins per season to flexibility alone, which is easily the difference between making and not making playoffs. Unlike the traditional process flexibility literature that focuses on the long chain, our results show that teams that derive more value from flexibility typically exhibit long subchains within the infield or outfield or both. This finding is primarily due to the fact that position players form two classes of players (infielder and outfielders) and it is less common to find a player who can play both types of positions at a high level of skill.
2. The least robust teams to worst-case injury are over four times as fragile as the most robust teams. Fragile teams tend to be built around one or two superstar players, which makes them susceptible to one or two targeted injuries. We also observe the classical trade-off found in the robust optimization literature between nominal performance and robustness, which in the baseball context is due to the limited supply of high-end talent.
3. Some teams can derive between a third and a half of their total value of flexibility from one or two players. Such players would be in the running for "Most Flexible Player", as defined by one of our new metrics. While individual flexibility always benefits the team, it may actually lead to reduced performance for the individual, which we observed for one quarter of players evaluated.
4. With targeted training of select players at select positions, it is possible to improve team robustness to worst-case injury by over 12%. However, the vast majority of potential player-position combinations that could be strengthened did not result in improved robustness, highlighting the value of our quantitative framework to pinpoint where effort should be concentrated to reduce fragility.
5. The coefficient of variation of the value a free agent adds to MLB teams can be over 40%, which indicates a player's value depends a lot on the specific team and roster he would be joining. Additionally, our framework can quantify the improvement in average and worst-case team performance due to a free agent signing, uncovering the trade-offs associated with acquiring different players and supporting data-driven decision making in roster construction.
6. Some teams can lose up to a quarter of their total production without their top player. For some players, their observed run production is close to what the team would lose without him. However, for others, their observable production may be an overestimate (upwards of 50%) of their true value, since those teams may have back-up flexibility that mitigates much of the potential loss. Our flexibility framework facilitates a rigorous "with or without you" analysis that can provide a more accurate valuation of a player for his team.

2. Literature

Baseball was the earliest sport to adopt analytics, driven in large part by the pioneering work of Bill James and his *Baseball Abstract* (James 1977). Many of the advanced baseball statistics that have been introduced since then are attributable to him. In addition to continued research into advanced statistics, modern operations research topics in baseball include scheduling (e.g., Trick et al. (2011)), in-game decision making (e.g., Ganeshapillai and Guttag (2014)), and drafting (e.g., Streib et al. (2011)). However, the topic of positional flexibility has not been explored previously in the literature, except for a short conference paper on which this paper builds (Chan and Fearing 2013).

Process flexibility has been studied extensively following the seminal work of Jordan and Graves (1995) on the impact of chaining and the value of the "long chain" in manufacturing and supply chains. Subsequent studies have extended the literature on flexibility to include multi-stage supply chains (Graves and Tomlin 2003, Hopp et al. 2010), queuing systems (Gurumurthi and Benjaafar 2004, Tsitsiklis and Xu 2012, Bassamboo et al. 2011), call centers (Wallace and Whitt 2005), workforce scheduling (Hopp et al. 2004), unbalanced networks (Deng and Shen 2013), assessing structural flexibility (Iravani et al. 2005), and analytical characterizations of the optimality of the long chain (Chou et al. 2010b, Simchi-Levi and Wei 2012). See Chou et al. (2008) for a comprehensive review.

Although most literature has considered manufacturing networks where production capabilities are binary (either a plant can or cannot produce a certain product), there has been some study of "continuous" product-specific manufacturing capabilities. For example, some papers assume that a plant should be highly capable when producing its primary product but less capable when producing secondary product. Such a reduction in capability may be realized in the form of increased production costs (Chou et al. 2010a, Mak and Shen 2009) or decreased productivity (Brusco and Johns 1998). With non-binary production capabilities, plants become non-identical in general. As a result, long chain structures may not be possible or at least degraded in value. This is what we see in baseball, where there tends to be less chaining between the infield and outfield positions.

Although most previous studies have focused on the value of flexibility in an expected value context, such as the increase in expected sales due to added flexibility, there has been recent and increasing interest in worst-case analysis in process flexibility (Chou et al. 2011, Simchi-Levi and Wei 2015, Wang and Zhang 2015, Yan et al. 2017)). Note that in a traditional supply chain context, both Tang and Tomlin (2008) and Simchi-Levi et al. (2015) have considered the use of flexibility as a means to improve robustness.

Robustness of flexible production networks has natural connections to vulnerability and interdiction. Attacker-defender (or max-min) models can be used by the attacker to devise optimal attack plans (see e.g., Alderson et al. (2011), Brown et al. (2006, 2009)), or to measure the worst-case impact of an attack from the perspective of the defender. We follow the latter interpretation to measure the impact of injuries on the performance of a baseball team.

3. Optimization models and metrics

Consider a baseball team that needs to choose from n different players (indexed by set J) to fill m different non-pitching, defensive positions (indexed by set I) over the course of a season, which is treated as a single period. In this paper, we ignore pitchers and focus on position players. The set J is based on the team's 40-man roster, thus including starters, bench players, and minor league depth options. For American League teams, we assume the designated hitter position is one of the defensive positions that needs to be filled – thus, $m = 9$ for American League teams and $m = 8$ for National League teams.

Let d_i be the total number of innings required for position i (i.e., demand for position i) and c_j be the total number of innings that player j is available (i.e., capacity of player j). Let x_{ij} be the number of innings that player j is assigned to position i during the season. We define v_{ij} to be the capability of player j at position i. The value of v_{ij} indicates how well-suited player j is at playing position i. We assume v_{ij} is the sum of both offensive and defensive contributions (see Appendix for more details). We also assume that the underlying bipartite graph describing the player-position combinations is complete, but that some edges (i, j) may simply have $v_{ij} = 0$. We will refer to the matrix $\mathbf{V}$ as the "capability matrix", which summarizes the relevant capability information for a team.

The nominal player-position assignment problem is:

$$
\begin{aligned}
\underset{\mathbf{x}}{\text{maximize}} \quad & \sum_{i \in I} \sum_{j \in J} v_{ij} x_{ij} \\
\text{subject to} \quad & \sum_{j \in J} x_{ij} \le d_i, \ i \in I, \\
& \sum_{i \in I} x_{ij} \le c_j, \ j \in J, \\
& x_{ij} \ge 0, \ i \in I, j \in J.
\end{aligned} \tag{1}
$$

Let $\mathbf{X}(\mathbf{c})$ denote the feasible region of formulation (1), parameterized by the vector of player capacities $\mathbf{c}$. The framework we develop, although rooted in the case of capacity uncertainty, is equally applicable to demand uncertainty given the equivalence in problem structure of formulation (1) between demand and supply. Restricting the matrix $\mathbf{V}$ to be binary recovers the standard production planning/network flow formulations considered previously in the literature (e.g., Chou et al. (2010b)).

We consider two approaches to modeling and solving the player-position assignment problem under injury-risk uncertainty, based on how the uncertainty is characterized. First, formulation (1) can be used directly when the uncertainty is described by a known probability distribution. In this case, we can simulate injuries and solve (1) using the simulated capacity vector $\mathbf{c}$ – this approach is essentially the approach of Jordan and Graves (1995) applied to capacity uncertainty. The optimal objective value of formulation (1) can then be written as

$$\mathbf{E}_{\mathbf{c}}\left[\max_{\mathbf{x}\in\mathbf{X}(\mathbf{c})}\sum_{i\in I}\sum_{j\in J}v_{ij}x_{ij}\right], \tag{2}$$

where the expectation is taken over the distribution of $\mathbf{c}$. To measure the value of flexibility in the expected value sense, two different capability matrices $\mathbf{V}^1$ and $\mathbf{V}^2$ can be compared based on their corresponding optimal objective values of (1) with the simulated capacity vector – we formalize this idea in Section 3.2.

The second approach we consider measures the robust (worst-case) value of flexibility by using a set-based characterization of the uncertain capacity vector. In this case, we assume nature is strategically choosing which players on the team to injure, within some injury budget, in order to maximize the degradation in team performance. The robust model is presented in the next section.

3.1. A robust player-position assignment model

To develop a robust optimization model, we first need to define the set of (uncertain) capacity vectors. Given nonnegative vectors $\mathbf{l},\mathbf{u},\boldsymbol{\sigma}$, we assume that the uncertain capacity vector lies in the following set:

$$\mathbf{C}_\Gamma=\left\{\mathbf{c}\in\mathbb{R}^n\,\middle|\,\sum_{j\in J}\frac{u_j-c_j}{\sigma_j}=\Gamma, l_j\le c_j\le u_j, j\in J\right\}. \tag{3}$$

The parameters l_j and u_j bound the extent to which the assumed capacity of player j can vary, while σ_j is a measure of spread within the corresponding interval $[l_j,u_j]$. We choose Γ to not exceed $\sum_{j\in J}(u_j-l_j)/\sigma_j$ so that $\mathbf{C}_\Gamma$ is non-empty. We use an equality instead of inequality ($\le$) in the first condition defining $\mathbf{C}_\Gamma$ because the worst-case will occur at equality for Γ in the range assumed; this definition of $\mathbf{C}_\Gamma$ helps with the subsequent model reformulation.

The structure of (3) is motivated by the limit theorem-inspired uncertainty sets described in Bandi and Bertsimas (2012). For capacity uncertainty, we believe it is most natural for capacity to decrease from its nominal value and not increase. Players can generally be expected to be available for the entire season unless an injury occurs. The parameter Γ is the *budget of uncertainty*. By varying Γ, we can study the value of flexibility across different levels of capacity uncertainty. Since our model measures the worst-case impact of uncertainty, the larger Γ is, the more damage nature can do.

As our focus is on quantifying the value of flexibility at a structural level, and not on developing a roster assignment that is robust to uncertainty, we formulate a min-max problem where the player-position assignment decisions are the inner decisions, and the player capacity realizations are associated with the outer minimization:

$$
\begin{aligned}
z_\Gamma = \min_{\mathbf{c}} \max_{\mathbf{x}} \quad & \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij} \\
\text{subject to} \quad & \sum_{j\in J} x_{ij} \le d_i,\ i\in I, \\
& \sum_{i\in I} x_{ij} \le c_j,\ j\in J, \\
& \sum_{j\in J} \frac{u_j - c_j}{\sigma_j} = \Gamma, \\
& l_j \le c_j \le u_j,\ j\in J, \\
& x_{ij} \ge 0,\ i\in I, j\in J.
\end{aligned} \tag{4}
$$

The optimal value of formulation (4) is a measure of the worst-case impact of capacity uncertainty. The nominal (no uncertainty) problem corresponds to $\Gamma = 0$ and has associated optimal objective value z_0. In the Appendix, we show that formulation (4) is equivalent to a tractable mixed-integer linear optimization model; this equivalent model is solved in Section 4 to generate the numerical results for the robust optimization experiments. For the typical problem size considered in Section 4 (~500 variables, ~1,000 constraints), the mixed-integer model solves in a few seconds using CPLEX 12.6. In the Appendix, we also provide an analytical characterization of the optimal $\mathbf{c}^*$ and $\mathbf{x}^*$ of model 4 for a special case, which gives insight into how nature chooses to injure players and how a team would assign players to positions.

3.2. Value of flexibility and robust protection levels

We introduce two metrics to measure the impact of flexibility on team performance. Of course, these metrics can be translated into the general production network context as well. Consider two capability matrices $\mathbf{V}^1$ and $\mathbf{V}^2$ such that $\mathbf{V}^2 \ge \mathbf{V}^1$ (component-wise). We think of $\mathbf{V}^2$ as being augmented from $\mathbf{V}^1$ with additional flexibility, or that $\mathbf{V}^1$ results from destroying some flexibility associated with $\mathbf{V}^2$.

The first metric, which we call the "average value of flexibility," considers the difference in expected performance by a team with capability matrix $\mathbf{V}^1$ versus $\mathbf{V}^2$ by simulating realizations of the uncertain capacity. This metric is applicable when there is an underlying probability distribution that describes the uncertain capacity. We define the average value of flexibility, $\text{AVF}(\mathbf{V}^1, \mathbf{V}^2)$, as

$$\text{AVF}(\mathbf{V}^1, \mathbf{V}^2) = \mathbf{E}_{\mathbf{c}}\left[\max_{\mathbf{x}\in\mathbf{X}(\mathbf{c})} \sum_{i\in I}\sum_{j\in J} v^2_{ij}x_{ij}\right] - \mathbf{E}_{\mathbf{c}}\left[\max_{\mathbf{x}\in\mathbf{X}(\mathbf{c})} \sum_{i\in I}\sum_{j\in J} v^1_{ij}x_{ij}\right]. \tag{5}$$

When normalizing AVF, we divide it by the term in the difference that represents the current team, depending on whether flexibility is being added (second term) or deleted (first term). In the former case, we measure how much value is gained if flexibility is added to the team and we denote the metric AVF^+. In the latter case, we measure the proportion of the overall performance of a team that is due to its flexibility and we denote the metric AVF^-.

The second metric, which we call the "robust value of flexibility," considers the difference in worst-case performance of a team with capability matrix $\mathbf{V}^1$ versus $\mathbf{V}^2$. We first define a team's "protection level" with respect to a relative loss of y (PL_y) as the largest value of Γ that guarantees at most such a loss in the worst-case – it provides a measure of the likelihood of not exceeding a specified loss level. For example, we define $\text{PL}_{0.1}$ as the largest value of Γ that results in at most a 10% loss in performance relative to the performance achieved when there is no capacity loss (i.e., the optimal value of the nominal problem when $c_j = u_j$ for all $j \in J$). A team that is more robust to injury would have a higher value for PL_y, given a loss level y. Formally, given a capability matrix $\mathbf{V}$, we define $\text{PL}_y(\mathbf{V})$ as

$$\text{PL}_y(\mathbf{V}) = \sup\left\{\Gamma \,\middle|\, 1 - \frac{\min_{\mathbf{c}\in\mathbf{C}_\Gamma} \max_{\mathbf{x}\in\mathbf{X}(\mathbf{c})} \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij}}{\max_{\mathbf{x}\in\mathbf{X}(\mathbf{u})} \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij}} \le y \right\}. \tag{6}$$

Using PL_y, the robust value of flexibility, $\text{RVF}_y(\mathbf{V}^1, \mathbf{V}^2)$, is defined as

$$\text{RVF}_y(\mathbf{V}^1, \mathbf{V}^2) = \text{PL}_y(\mathbf{V}^2) - \text{PL}_y(\mathbf{V}^1). \tag{7}$$

Roughly speaking, $\text{RVF}_y(\mathbf{V}^1, \mathbf{V}^2)$ measures how many more standard deviations of uncertainty a team with $\mathbf{V}^2$ is protected against over a team with $\mathbf{V}^1$, at a certain loss level y. Similar to AVF, we derive a relative measure by dividing by $\text{PL}_y(\mathbf{V}^1)$ or $\text{PL}_y(\mathbf{V}^2)$, depending on whether we wish to measure the value of additional (RVF^+) or existing (RVF^-) flexibility, respectively.

4. Baseball case studies

In this section, we demonstrate the application of our models and metrics using data from Major League Baseball. The six subsections each focus on a different set of computational experiments that aim to answer the six main questions posed in the Introduction: 1) What is the average value of flexibility at a team level? 2) How do we measure the robustness of a team to worst-case injury? 3) What is the average value of flexibility at the individual level? 4) Is it possible to improve team robustness to worst-case injury using only the existing roster? 5) Given a list of free agents, which one should a team acquire to improve its performance the most? 6) Who is the Most Valuable Player?

To perform our studies, we utilize MLB 40-man rosters prior to the start of the 2012 season, including all players who started the season on the disabled list. Our study is a prospective analysis of team composition, so we use data on projected player performance over the 2012 season. For offensive performance, we use the well-known ZiPS offensive projection method (Szymborski 2012), which we convert into runs per game. For defensive performance, we develop a statistical model that simultaneously estimates defensive capabilities at multiple positions using a linear regression with correlated random effects. Of particular note is that the assumed multivariate-normal covariance structure allows us to estimate a player's capabilities at each position, including positions he has previously not played. As with offensive performance, the defensive numbers are converted into runs per game (runs saved, in this case). We assume that offensive production is independent of the defensive position a player is assigned. Thus, differences in performance between a player playing two different positions is assumed to be entirely due to the difference in ability to prevent runs in those two positions.

Each player's total production per game is measured relative to the expected contribution of an off-roster, readily-available, minimum-cost alternative, also known as a *replacement-level* player (FanGraphs 2014). Thus, a player with a production rate of 0.25 runs per game would be expected to create 0.25 runs per game of incremental value over a replacement player, through either runs scored or saved – such runs are called *runs above replacement* (RAR). RAR is the measure of productivity we use in the experiments below (v_{ij} is the per-game RAR of player j at position i). Replacement level is an important concept since it provides a baseline against which to measure performance and the value of flexibility in the face of injury. We note that trades for players better than replacement level are one way to address injuries mid-season. However, such trades are uncommon. Relatively few trades occur during the season overall and most occur during the run-up to the trade deadline between teams on opposite sides of the playoff cutoff. Thus, while our approach may overestimate the impact of poor flexibility in dealing with injuries, the overestimation is likely to be small.

For the range of winning percentages in which most teams fall (40% to 60%) the relationship between runs scored/saved and wins is approximately linear with ten runs corresponding to a win (FanGraphs 2014). Thus, a player who creates 0.25 runs per game would be worth 40.5 runs over a full season or approximately four additional wins. As opposed to other team sports, player contributions in baseball can be considered relatively independent, especially offensive contributions. There are many approaches to modeling overall team performance based on individual player performance, with additive models constituting a major class (Wyers 2010). So for tractability in our model, we approximate overall team production as the sum of individual contributions. Throughout, we treat the demand for each position as deterministic and equal to 1,458 innings (nine innings per game times 162 games per season).

To estimate injury risks, we develop a two-stage regression model: a logistic regression using age as the independent variable to estimate the probability of a serious injury resulting in a trip to the disabled list; and a log-linear regression to estimate the duration of the injury. In addition to accurate estimates of mean and variance, the two-stage model provides a good fit to the shape of the historical injury distribution.

The high level approach to calculating the value of flexibility and robust protection level metrics – considering all data sources, statistical models and optimization models – is outlined in Figure 1. A complete description of the data and statistical models for defensive performance and injury are available in the Appendix.

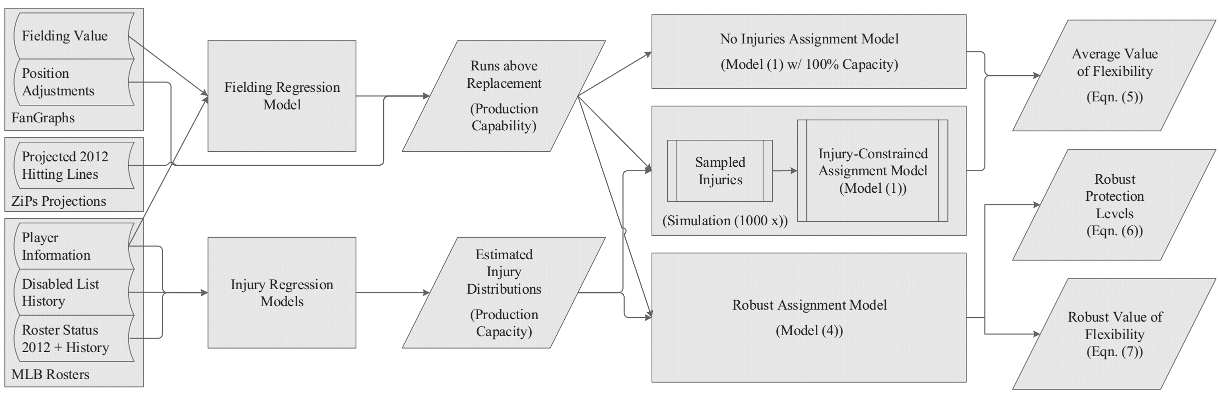

Figure 1 **Data and models used to calculate flexibility metrics.**

4.1. Value of flexibility at a team level

First, we compare team performance subject to simulated injuries when the team's players have their native, non-zero capabilities at their secondary positions (the *flex* case) to the case when the team's players are only generating value at the primary assigned position (the *no-flex* case). To determine a player's primary assigned position, we solve model (1) without considering injuries and record the position to which each player is assigned. The team's original capability matrix corresponds to the *flex* case, while a modified version of the matrix where the capabilities of the non-primary positions are assigned a value of 0 is used to represent the *no-flex* case. By virtue of the way we defined the *no-flex* case, injuries to the starters result in replacement-level players playing at those positions. In other words, we assume no roster depth beyond replacement level.

To calculate the value of flexibility at a team level, we first simulate injuries by sampling player capacities based on the injury distributions. These simulated values are used as input to the nominal model (1), which is then solved to determine the value-maximizing assignment of players to positions over the course of a 162-game season. We repeat this process 1000 times for both the *no-flex* and *flex* capability matrices ($\mathbf{V}^1$ and $\mathbf{V}^2$, respectively) and calculate $\text{AVF}^-(\mathbf{V}^1, \mathbf{V}^2)$.

Figure 2 summarizes the average value of flexibility for each team. The value that flexibility contributes to team performance varies substantially across the league, ranging from 4.3% for the Baltimore Orioles to 12.0% for the Chicago Cubs. For each team, we can convert the value of flexibility metric into runs-generated, which we use to approximate the number of wins a team will earn due solely to flexibility. Assuming that 10 runs is equivalent to one win, flexibility alone generates between 0.7 to 2.7 wins for each team. Furthermore, the top 10 teams in the league can attribute at least 1.5 wins due to flexibility alone. A couple wins is significant because it is typically the difference between making and missing the playoffs.

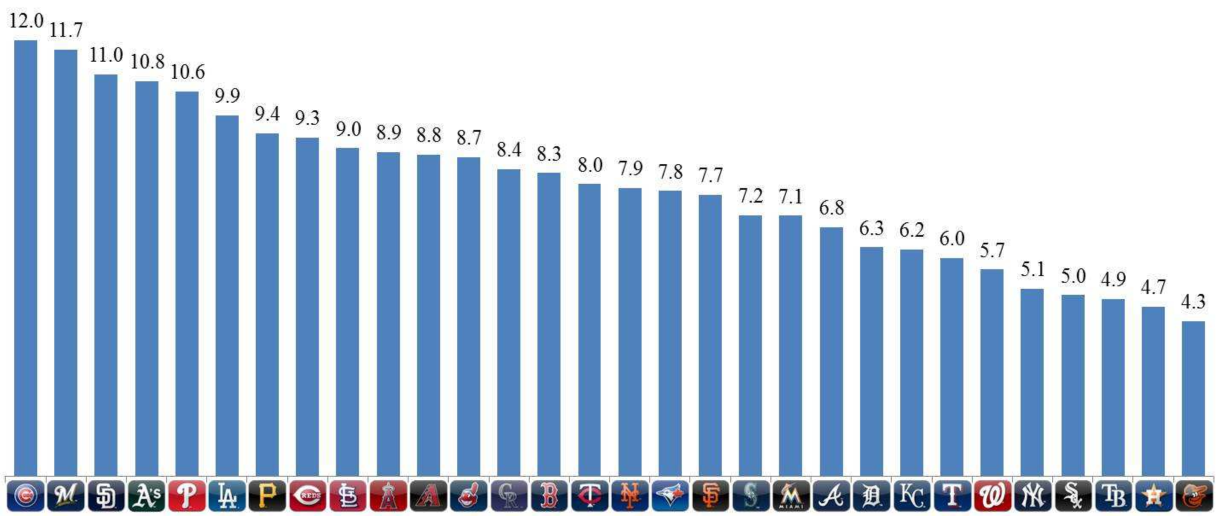

Figure 2 AVF$^-$ for all thirty MLB teams (%). (See Table 11 for correspondence between logos and team names.)

Next, we examine a few teams more closely to understand how flexibility is creating value. Tables 2, 3 and 4 illustrate how flexibility impacts the Chicago Cubs, Milwaukee Brewers and Baltimore Orioles, respectively. The tables can be read as follows. The primary "matrix" lists the non-zero v_{ij} values, representing the per-game runs above replacement for each player at each position. Shaded cells indicate the playing time of a player at the specified position in the *flex* case with the grayscale proportional to playing time, ranging from the entire season (black) to no games (white). The total amount of playing time for each player is given in the "Util." column as a percentage of the total number of innings in the season (1,458). The starters are all assigned their maximum available playing time, which was determined through simulation. The "RAR" column indicates how many runs above replacement are generated by the player over the season, based on their v_{ij} values and playing time at the assigned positions.

Table 2 shows that both roster depth and positional flexibility create value in Chicago. There is chaining in the infield positions through bench players such as Bryan LaHair (who fills in at 1B and 3B when the starters are injured), Darwin Barney (at 2B, SS) and Luis Valbuena (at 2B, SS) and chaining in the outfield positions through Alfonso Soriano (at LF), Tony Campana (at LF, CF, RF), and Reed Johnson (at LF, CF). Milwaukee (Table 3), the team that creates the second most value from flexibility has similar infield and outfield chains, but exhibits another interesting trait – a single player who connects the infield and outfield chains. The primary shortstop Mat Gamel, who can play many positions with skill, is the next best option at 1B, 3B, and LF after the respective starters. Thus, he helps Milwaukee move towards a long chain structure by connecting the infield and outfield. The reason he has capacity to fill in as needed in the other positions is because Alex Gonzalez is a capable back-up shortstop. Indeed, Gonzalez plays almost half the season at SS with only a small decrease in v_{ij} (from 0.14 to 0.11) enabling Gamel to spend time at the other positions as needed. However, it should be noted that a high-capability link such as Gamel between the infield and outfield, coupled with suitable back-up capacity, is rare; almost all of the top-ranked teams in AVF$^-$ exhibited long subchains in the infield or outfield, rather than a long chain spanning the infield and outfield. This finding is likely characteristic of two things: 1) Infield and outfield positions are very different. Players typically specialize into only one of these two classes, and often only a specific position within a class, in the case of inflexible players. 2) bench players rarely perform at a comparable level as starting players.

Table 4 shows that Baltimore's ability to create an infield chain is limited (Baltimore is the last-ranked team in AVF$^-$, cf. Figure 2), and not without serious degradation in performance – Robert Andino provides a modest amount of infield flexibility, but has substantially lower capability values than the starters at 2B, 3B and SS. There is even less flexibility in the outfield. Although starting infielders Nick Markakis, Brian Roberts and Mark Reynolds have comparable outfield capabilities as the starting outfielders, and therefore could help connect the infield and outfield chains, there is minimal excess infield capacity with enough capability above replacement level performance to take advantage of their flexibility.

We note that sometimes the player-position assignment produced by the model does not reflect the typical real-world assignment of a player. Our player-position assignments are determined from the optimal solution of a model with projected v_{ij} values. Both the method of determining the assignment (optimization) and the estimation of the v_{ij} values may contribute to any possible mis-matches between our assignments and reality. Furthermore, managers may be making suboptimal decisions or have different valuations for v_{ij} than what our projections suggest. A concrete example is Nick Markakis, who is discussed above. In reality, he is the starting RF for Baltimore but our model places him at 1B (a position he has experience playing). Chris Davis is the actual 1B starter

Table 2 Player-position capabilities and playing time distribution in AVF$^-$ calculation for the Chicago Cubs

	v_{ij} values								Util. (%)		RAR	
Name	C	1B	2B	3B	SS	LF	CF	RF	*no-flex*	*flex*	*no-flex*	*flex*
Geovany Soto	0.25								82	82	33.6	33.6
Anthony Rizzo		0.09				0.07		0.05	85	85	12.9	12.9
Blake DeWitt			0.12	0.13	0.08	0.07		0.01	84	84	16.9	17.0
Ian Stewart			0.12	0.14	0.08	0.05		0.06	83	83	18.1	17.7
Starlin Castro			0.14	0.12	0.21				85	85	28.7	28.7
David DeJesus						0.19	0.14	0.14	82	82	25.0	24.6
Dave Sappelt						0.11	0.12	0.05	84	84	16.0	16.0
Marlon Byrd						0.11	0.12	0.12	82	82	16.1	16.1
Steve Clevenger	0.14									16		3.8
Welington Castillo	0.13									2		0.3
Bryan LaHair		0.07		0.11		0.06		0.06		38		5.9
Darwin Barney			0.07	0.07	0.09					21		2.7
Luis Valbuena			0.06	0.02	0.06	0.04				3		0.3
Alfonso Soriano						0.15	0.04			20		4.7
Tony Campana						0.12	0.11	0.11		30		5.5
Reed Johnson						0.10	0.01			1		0.1

Shading is proportional to playing time in the *flex* case: a darker cell means more playing time at that position. The players along the main diagonal are the starters in the *no-flex* case. Util. = fraction of season the player is playing

Table 3 Player-position capabilities and playing time distribution in AVF$^-$ calculation for the Milwaukee Brewers

	v_{ij} values								Util. (%)		RAR	
Name	C	1B	2B	3B	SS	LF	CF	RF	*no-flex*	*flex*	*no-flex*	*flex*
George Kottaras	0.18								83	83	24.4	24.4
Corey Hart		0.11				0.12	0.13	0.15	83	83	15.3	15.3
Rickie Weeks		0.12	0.25	0.17	0.20	0.16		0.13	83	83	34.1	34.0
Aramis Ramirez		0.11	0.18	0.24	0.19	0.14		0.12	81	81	30.8	30.8
Mat Gamel		0.09	0.13	0.17	0.14	0.18		0.10	85	85	19.1	19.4
Ryan Braun		0.19				0.28	0.20	0.15	83	83	37.2	36.8
Norichika Aoki						0.06	0.13	0.06	83	83	18.1	18.1
Caleb Gindl						0.08	0.07	0.11	83	83	15.3	15.3
Jonathan Lucroy	0.14									16		3.6
Martin Maldonado	0.04									1		0.1
Taylor Green			0.14	0.13	0.11	0.08		0.08		27		5.8
Alex Gonzalez			0.01	0.02	0.11					49		8.5
Carlos Gomez						0.11	0.13	0.11		28		5.7
Nyjer Morgan						0.11	0.12	0.11		12		2.1

Shading is proportional to playing time in the *flex* case: a darker cell means more playing time at that position. The players along the main diagonal are the starters in the *no-flex* case. Util. = fraction of season the player is playing

but our model puts him at backup at that position. This result is likely due to the process of determining the assignment through optimization as opposed to estimation error since Markakis'

Table 4 Player-position capabilities and playing time distribution in AVF$^-$ calculation for the Baltimore Orioles

	v_{ij} values									Util. (%)		RAR	
Name	C	1B	2B	3B	SS	LF	CF	RF	DH	*no-flex*	*flex*	*no-flex*	*flex*
Matt Wieters	0.25	0.06							0.07	83	83	34.1	34.1
Nick Markakis		0.12				0.10	0.08	0.13	0.09	84	84	16.2	16.5
Brian Roberts		0.04	0.20	0.13	0.14	0.05		0.07	0.05	83	83	26.8	26.8
Mark Reynolds		0.08	0.12	0.15	0.12	0.12		0.08	0.09	82	82	20.7	20.0
J.J. Hardy			0.16	0.18	0.27				0.05	83	83	35.6	35.6
Wilson Betemit		0.01	0.09	0.08	0.09	0.10				81	81	13.0	13.0
Adam Jones						0.10	0.17	0.13	0.07	85	85	23.1	23.1
Nolan Reimold		0.03				0.07	0.04	0.08	0.04	83	83	10.3	10.3
Chris Davis		0.03		0.05		0.02			0.02	84	84	2.7	3.1
Ronny Paulino	0.07										15		1.7
Robert Andino			0.05	0.06	0.06						42		4.2
Ryan Flaherty			0.06								15		1.4
Josh Bell				0.05							5		0.4
Endy Chavez						0.02	<0.01	0.01			21		0.3

Shading is proportional to playing time in the *flex* case: a darker cell means more playing time at that position. The players along the main diagonal are the starters in the *no-flex* case. Util. = fraction of season the player is playing

v_{ij} at RF is greater than his v_{ij} at 1B, and since Nolan Reimold is a better second option at RF than Davis is at 1B.

Overall, our results reinforce the importance of not only having flexible players, but having bench players who are skilled enough to not seriously degrade team performance compared to the starters. The value that flexibility creates for a team varies widely across the league and depends as much on the flexibility of the starters as it does on the capability of the bench players. Our results highlight some interesting differences from the traditional process flexibility literature. In particular, the teams that exhibited the most flexibility generally did not have a long chain, but instead had two long subchains, one in the infield and one in the outfield. This result is not surprising given that players tend to specialize as infielders or outfielders but not both. Additionally, bench players often represent a sizeable drop-off in performance compared to starting players. Thus, even if a team has a flexible player who can connect the infield and outfield, it may not be able to capitalize on that flexibility without suitable bench players.

Although not explicitly studied in this paper, we note that flexibility creates value in other ways. First, flexibility allows a team to create *platoon* advantages, which arise when a team's batter faces a pitcher of opposing handedness (e.g., a left-handed batter typically performs better against a right-handed pitcher). There is a roughly 70/30 split between right and left-handed pitchers in the league, so knowing the handedness of the players who are assigned playing time means flexibility creates value even in the absence of injuries. Second, flexibility gives a team additional options for in-game strategy, such as enabling defensive adjustments where better defenders are substituted into the game, which is common in the late stages of close games. Third, no player is immune to the risk of poor performance, and flexibility again gives a team options to mitigate this risk. Platooning has been explored in a related paper (Chan and Fearing 2013), while in-game strategy and poor performance risk is very context dependent.

4.2. Robust protection levels at a team level

Next, we quantify the benefit of flexibility in the face of worst-case injury outcomes. We compute $\mathrm{PL}_{0.1}(\mathbf{V})$ values for each team in the league, given a capability matrix exhibiting flexibility. The PL values were determined by iteratively solving the robust model (14) shown in the Appendix, increasing Γ by 0.25 at each step, and linearly interpolating between the Γ values that bound the specified loss. The parameters of the uncertainty set $\mathbf{C}_\Gamma$ were set as follows: the maximum capacity value assumed no injuries ($u_j = 1$); the minimum capacity value, l_j, was set to one minus the 90th percentile of the injury distribution; the standard deviation, σ_j, was calculated directly from the injury distribution. The injury distribution depends on the player's age, which means that both l_j and σ_j vary by player.

Figure 3 summarizes the $\mathrm{PL}_{0.1}$ values for each team. The PL metric identifies which teams are the least (high PL value) and most (low PL value) susceptible to worst-case injury. From the figure, we see that the least robust team (Washington) is over four times as fragile as the most robust team (Oakland). That is, to reduce the team's performance by 10%, nature requires a budget Γ for Oakland that is over four times as large as the budget for Washington.

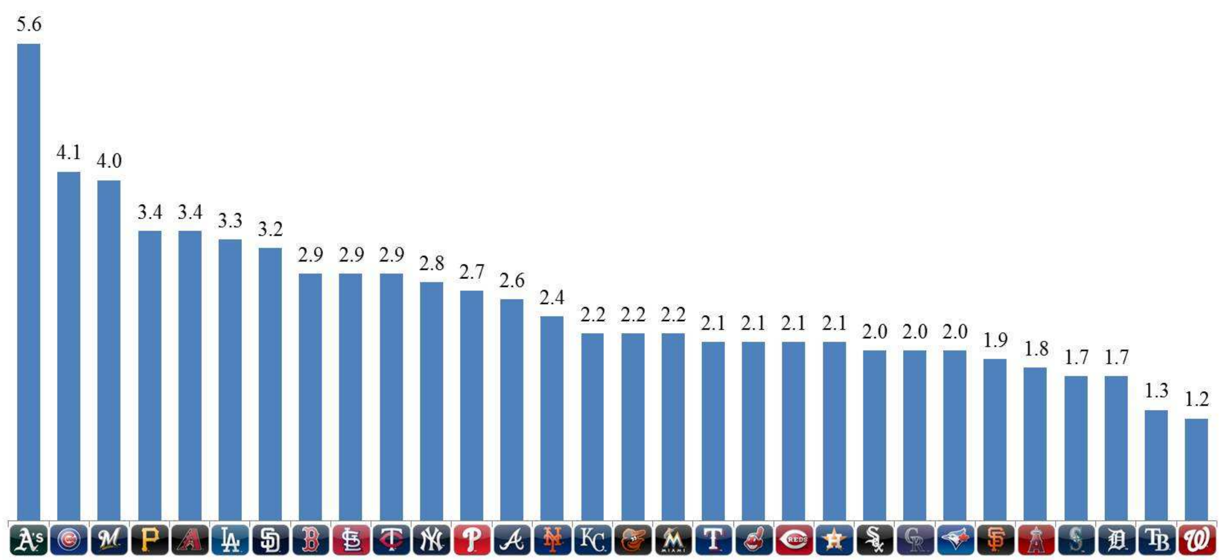

Figure 3 $\mathrm{PL}_{0.1}$ **for all thirty MLB teams. (See Table 11 for correspondence between logos and team names.)**

Next, we take a closer look at the rosters for Washington and Oakland. Table 5 shows that the low PL value for Washington is due to a large drop-off in capability at the 3B position after their superstar third baseman Ryan Zimmerman. We see that Ryan Zimmerman is the only player injured, first baseman Michael Morse fills in for Zimmerman at 3B at a substantial decrease in capability (0.10 vs. 0.35), and a bench player plays 1B when Morse is at 3B, representing a further drop in performance at 1B (0.03 vs. 0.12). The net reduction in capability value is 0.34 through the chained losses at 3B and 1B. Note that chaining can be easily visualized in this case (see arrows in Table 5): starting from the diagonal entry corresponding to Zimmerman, move up to the shaded cell representing who plays 3B when Zimmerman is injured (Morse), move left to Morse's primary position of 1B, and the move down to the shaded cell representing who plays 1B when Morse is playing 3B (LaRoche). Other teams with a low PL value are similar in that they have a spike in capability in one or two positions, which leads to severely degraded team performance from a single injury or two.

Table 5 Player-position capabilities and playing time distribution in $\mathrm{PL}_{0.1}$ calculation for the Washington Nationals

	v_{ij} values								
Name	C	1B	2B	3B	SS	LF	CF	RF	Util. (%)
Wilson Ramos	0.25								100
Michael Morse		0.12		0.10		0.10		0.08	100
Stephen Lombardozzi			0.11	0.01	0.03	0.01		0.01	100
Ryan Zimmerman		0.17	0.26	0.35	0.27	0.18		0.23	61
Danny Espinosa			0.11	0.05	0.10	0.03		0.04	100
Jayson Werth		0.08				0.18	0.19	0.18	100
Brett Carroll						0.06	0.09	0.13	100
Bryce Harper						0.07	0.09	0.13	100
Jesus Flores	0.04								0
Jhonatan Solano	0.03								0
Adam LaRoche		0.03							39
Ian Desmond			0.03		0.06				0
Roger Bernadina						0.07	0.03	0.01	0
Rick Ankiel						0.03	0.03	<0.01	0

Shading is proportional to playing time for $\Gamma = 1.25$: a darker cell means more playing time at that position. The players along the main diagonal are the starters when $\Gamma = 0$. Util. = fraction of season the player is playing. Arrows illustrate player-position chaining as a result of injury.

Oakland's roster, on the other hand, has a much more complex and balanced response to injuries, as seen in Table 6. Scott Sizemore's playing time at 3B is reduced via injury and made up for by the infield chain involving Daric Barton and Kila Ka'aihue. Although injuring Sizemore does not result in an immediate drop in capability at 3B, the resulting chain (see arrows in table) induces a drop in capability at 1B. In the outfield, Josh Reddick gets injured at RF and is replaced by a replacement level (zero value) player. Another chain is formed (see arrows in table) when Yoenis

Cespedes gets injured at CF: Coco Crisp swings over from LF to cover CF most of the time but he also gets injured, drawing Seth Smith over to cover LF, bench player Michael Taylor to cover CF, and leaving replacement level players to satisfy the remaining demand at DH. The chaining in the outfield to reduce the impact of injuries is made possible by the comparable capability of the three outfielders and the designated hitter.

Table 6 Player-position capabilities and playing time distribution in $\mathrm{PL}_{0.1}$ calculation for the Oakland Athletics

	v_{ij} values									
Name	C	1B	2B	3B	SS	LF	CF	RF	DH	Util. (%)
Kurt Suzuki	0.17									100
Daric Barton		0.07		0.08				0.02	0.01	100
Wes Timmons			0.10	<0.01	0.02					100
Scott Sizemore			0.09	0.08	0.04	0.03		0.01		73
Cliff Pennington			0.07	0.05	0.10					100
Coco Crisp						0.13	0.15	0.09	0.02	50
Yoenis Cespedes						0.11	0.18	0.11	0.05	52
Josh Reddick						0.07	0.08	0.07		53
Seth Smith		<0.01				0.11	0.07	0.05	0.02	100
Derek Norris	0.14									0
Josh Donaldson	0.08									0
Kila Ka'aihue		0.02		0.02						27
Chris Carter		<0.01		<0.01						0
Jemile Weeks			0.08	0.01	0.03	<0.01				0
Eric Sogard			0.06	0.01	0.07					0
Collin Cowgill						0.08	0.05			0
Michael Taylor						0.02	0.09	0.02		45

Shading is proportional to playing time for $\Gamma = 5.5$: a darker cell means more playing time at that position. The players along the main diagonal are the starters when $\Gamma = 0$. Util. = fraction of season the player is playing. Arrows illustrate player-position chaining as a result of injury.

There are also some notable differences in rank observed between Figure 2 (AVF^-) and Figure 3 ($\mathrm{PL}_{0.1}$). Some teams experienced large shifts in their rankings, such as the Los Angeles Angels (10th in AVF^- and 26th in $\mathrm{PL}_{0.1}$) and the Baltimore Orioles (30th in AVF^- and 16th in $\mathrm{PL}_{0.1}$). Referring back to Table 4, Baltimore does not have a single player that creates a substantial imbalance in the team. So even though Baltimore ranks poorly in terms of AVF^-, which is due to a weak bench, the team is not overly susceptible to worst-case injuries. On the other hand, the Los Angeles Angels' star infielder Albert Pujols is the cause of his team's fragility because he is so much better than his teammates. However, the Angels also have a deep and capable bench, so their expected performance with respect to injury risk is above average.

Finally, we comment on the relationship between nominal performance and robustness, as measured by z_0 and $\mathrm{PL}_{0.1}$, respectively (Figure 4). We see that the most robust teams tend to have among the lowest levels of nominal performance, which is consistent with the fact that those teams tend to lack superstars. Without superstars, they generally perform below average, but also are insensitive to selective injuries by nature. At the other end of the spectrum, there are many teams that are high-performing but fragile, which suggests that one or two superstars are driving a lot of performance for that team. There are also teams that do poorly in both metrics. These teams have some players that stand out from their teammates, leading to increased fragility, but their performance is still generally poor, resulting in poor team performance overall. Interestingly, but perhaps not surprisingly, there do not appear to any be teams that excel at both nominal performance and robustness. If we examine the five teams on the efficient frontier generated by these two objectives (the white markers in Figure 4) and run a linear regression ($R^2 = 0.84$), we find that for every unit improvement in $PL_{0.1}$, team performance is reduced by 68.4 runs, which is almost seven losses, a significant amount. To do well in both dimensions likely requires several superstars who have capable back-ups and flexibility distributed among the starting and bench players, which does not appear possible in very competitive sports with limited high-end talent pools.

4.3. Value of individual flexibility: Who deserves the "Most Flexible Player" award?

Replicating the framework used to calculate the value of flexibility at the team level, we can also determine the value of a player's individual flexibility to his team by comparing team performance when a player has his native secondary capabilities against the situation when the same player has replacement level (i.e., zero) secondary capabilities, leaving all other players unchanged. In this experiment, we consider a player's primary position to be the one in which he spends the plurality of his time when solving the nominal problem with simulated injuries. This method of determining a player's primary position is different from Section 4.1 because we want to ensure bench players are included in the individual flexibility measurement, whereas the team flexibility experiment focused on the value of flexibility of the starting lineup as a whole.

Table 7 highlights the top 10 players who create the most team value from their individual flexibility. We see that sometimes a team derives much of its total value of flexibility from only a couple players on the roster. For example, Brandon Belt, who is first on the list, is responsible for almost one third of San Francisco's value of flexibility (cf. Figure 2). Two players from Texas are in the top 10 – Michael Young and Nelson Cruz – and are responsible for more than half of their team's value of flexibility. Overall, the metric $\mathrm{AVF}^{-}(\mathbf{V}^1, \mathbf{V}^2)$, where $\mathbf{V}^1$ and $\mathbf{V}^2$ differ by the flexibility of a single player, defines what we refer to as the "Most Flexible Player". That is, players who rank highly in this metric are those whose flexibility contributes the most to their team's performance. Equivalently, a team who loses this player's flexibility stands to lose the most in team performance.

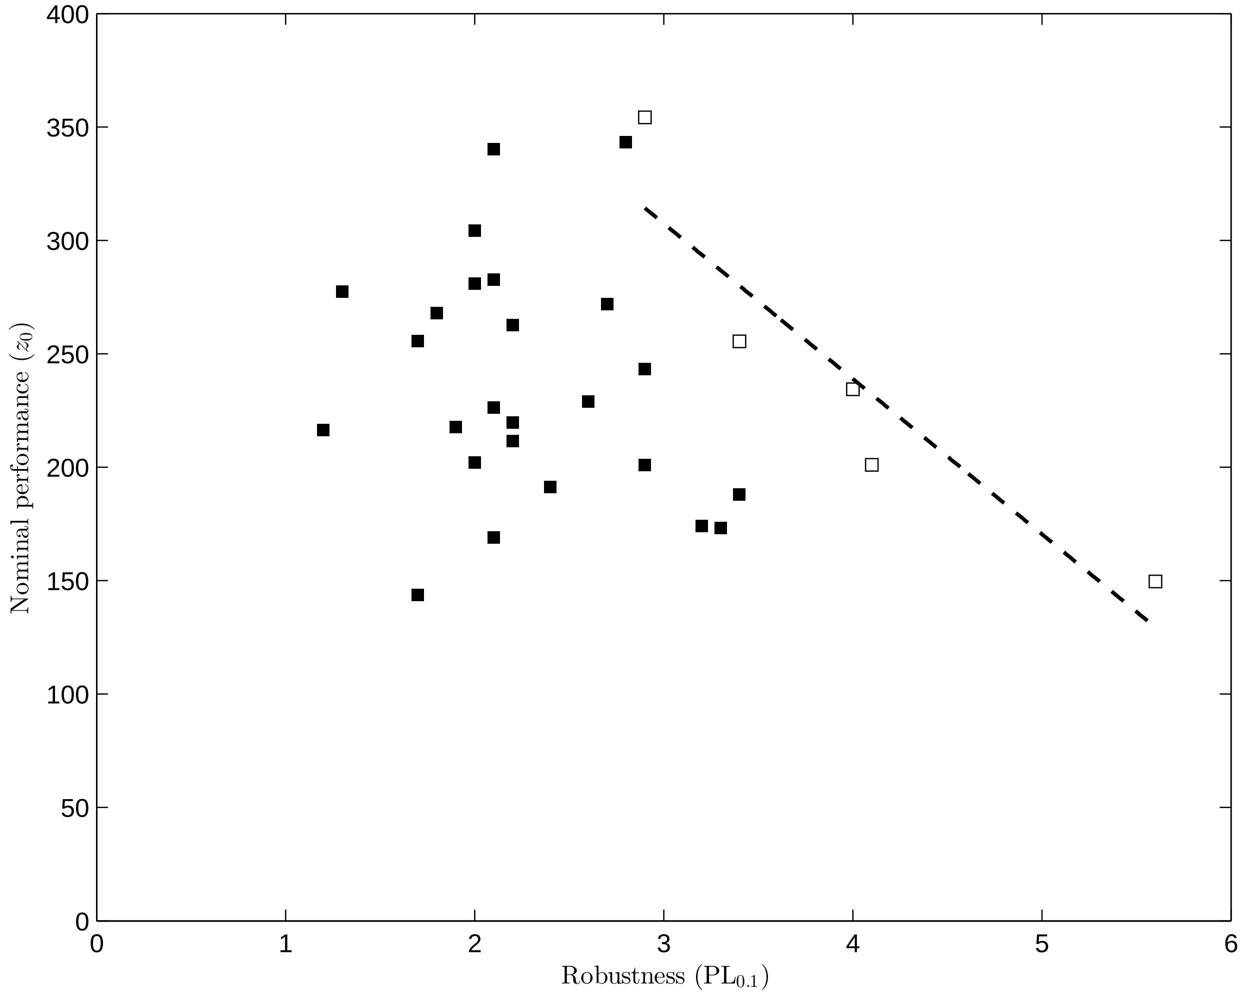

Figure 4 **The relationship between nominal performance in RAR (optimal value of formulation (1)) and robustness (10% protection level)**

Table 7 **Top 10 players whose individual flexibility creates the most value (AVF^-) for their team. (See Table 11 for correspondence between team names and abbreviations.)**

Player	Team	AVF^- (%)
Brandon Belt	SF	2.49
Michael Young	TEX	2.12
Nick Swisher	NYY	1.69
Victor Martinez	DET	1.41
Kyle Seager	SEA	1.27
Torii Hunter	LAA	1.26
Nelson Cruz	TEX	1.22
Mat Gamel	MIL	1.17
Jayson Werth	WSH	1.12
Placido Polanco	PHI	0.95

4.3.1. The downside of flexibility While it is intuitive that individual flexibility benefits the team, next we examine whether there is a potential downside to flexibility. Namely, is it possible that an individual's flexibility can hurt his own production? Figure 5 examines the impact of an individual player's flexibility on both his performance (x-axis) and his team's performance (y-axis) in terms of runs above replacement. First off, we see that all values in the y-direction are positive, which confirms that flexibility always benefits the team. However, a player's individual production may decrease as a result of his flexibility. For example, Jose Bautista's flexibility costs him about four runs, while his team gains two runs. Recall that a player's production is the sum of (position-independent) offensive and (position-dependent) defensive contributions, so a decrease in individual production is due to a decrease in runs saved by playing a different defensive position. Bautista is an all-star player for Toronto (in 2012) who can play multiple positions with skill. As a result, he may be asked to play secondary or tertiary positions, which he plays at a lower level of capability compared to his primary position, when a teammate is injured. By moving away from his primary position, he enables a chain that mitigates the loss in team production due to the injured player, at the expense of his own production.

On the other hand, Brandon Belt (who creates the most team value through individual flexibility, cf. Table 7) is an example of a player whose flexibility also provides substantial benefit to his own production. In his nominal assignment, he is sacrificing a bit of individual production for the benefit of his team. However, his flexibility allows the team to play him at a more productive position in the case of injury. The Bautista and Belt examples illustrate how flexibility impacts a player's individual performance depending on when a player is assigned a suboptimal position: either after an injury occurs (Bautista) or before an injury occurs (Belt).

Out of 489 players, individual flexibility results in lower production for 112 players, higher production for 143 players, and unchanged production for 234 players. Conditional on a player's run production changing due to their flexibility, 44% (112/255) of players saw a decrease in production while 56% (143/255) saw an increase. For players whose production decreased, the average loss was about half a run (roughly 3%) and the bottom 10% saw their production decrease by over 12%. For players whose production increased, the average gain was about seven tenths of a run (roughly 17%) and the top 10% saw their production increase by 36%. A summary is provided in Table 8.

To confirm the uniqueness of our solutions, we spot checked the optimal solutions. This check is important to rule out the possibility that there is a multiplicity of optimal solutions and only some of them may exhibit the observed "downside of individual flexibility" phenomenon. We believe uniqueness holds throughout since our capability values come from a continuous distribution over a population, and thus the ability to re-distribute capacity on a costless basis across different players is very unlikely, essentially measure zero. However, in a more stylized setup with unit capabilities, which is the typical setup in most manufacturing flexibility studies with "on-off" flexibility, or even

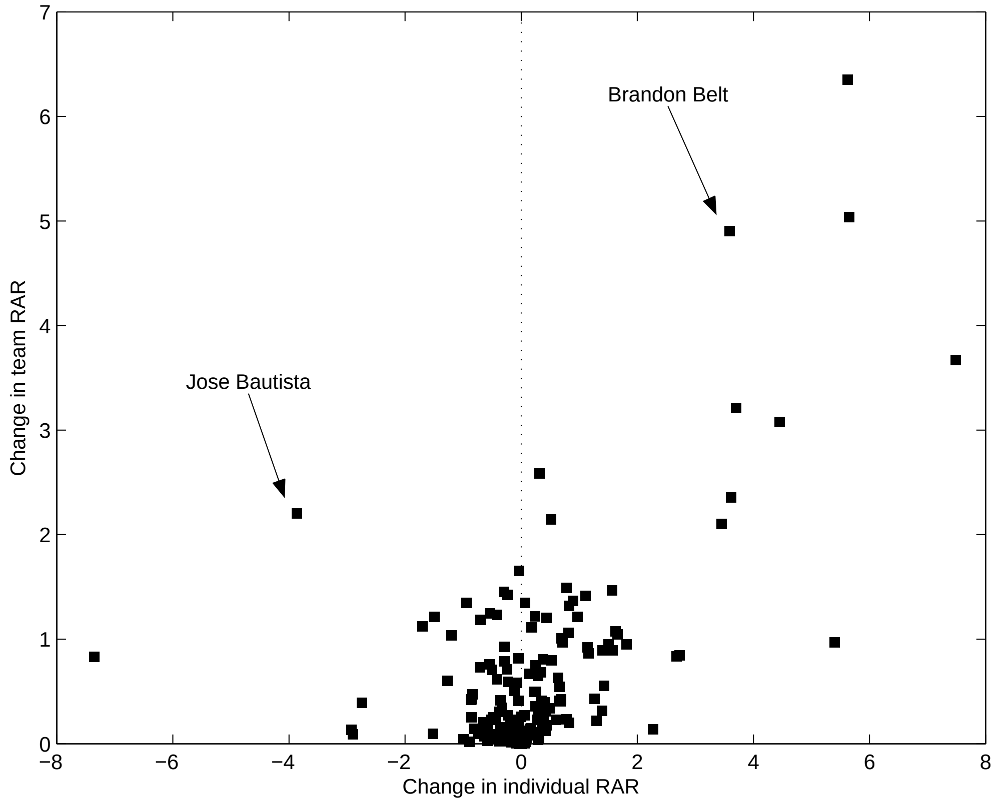

Figure 5 **The impact of individual flexibility on both team and individual performance. RAR = runs above replacement.**

Table 8 **Summary of individual flexibility on individual and team performance**

Flexibility impact on player	Number of players	Production change	
		Individual	Team
Lose production	112	−0.47 (−3.1%)	0.30 (0.14%)
Gain production	143	0.67 (16.6%)	0.56 (0.26%)
Lose or gain	255 (52%)	0.17 (7.9%)	0.45 (0.2%)
No effect	234 (48%)	-	-
Total	489	0.09 (4.1%)	0.23 (0.0%)

Values in the “Individual” and “Team” columns represent the average gain/loss in individual and team runs, respectively.

capabilities that take on values from a discrete set, the likelihood of multiple optimal solutions may be non-trivial.

Overall, individual flexibility may be an advantage or disadvantage for a player. A player who is flexible and can fill in capably at a secondary position may contribute substantial value to his team, but risk sacrificing his own production if the manager plays him at (individually) suboptimal positions. In practice, players are often hesitant to move away from their primary position because the market tends to value individual production. For example, contracts and performance bonuses are currently predominantly based on individual production. In principle, a player’s contribution to his team's overall performance should form part of a player's formal contract evaluation in order to better align player and team incentives. Our value of flexibility approach provides one possible way to do so.

4.4. Improving team robustness by developing secondary capabilities

In this section, we examine how a small increase in a player's secondary capability can impact his team's robustness to worst-case injury using the RVF$^+$ metric. Such an improvement in a player's secondary capability could in principle be achieved through vigorous training before the start of the season. The goal of these computations is to show that a manager can potentially use our framework to identify targeted opportunities to improve worst-case performance even without any changes to the roster. In our experiments, we add two runs per 150 games to all the secondary v_{ij} values of each player who did not primarily play DH historically, who played less than 45 innings at the secondary position, and whose secondary position was not their nominal assignment. The addition of two runs per 150 games represents a fairly small increase in their secondary capability – enough to realize a non-trivial increase in network robustness, but not so large as to be unreasonable for the short training period. This increase is captured in a capability matrix $\mathbf{V}^2$, which has a slight increase to a single component of the base matrix $\mathbf{V}^1$.

Out of the 670 player-position capability augmentations considered, 45 of them resulted in an increase in his team's robustness to worst-case injury, as measured by RVF$^+$. The top 10 player-position capability augmentations (in terms of increase in RVF$^+$) are shown in Table 9 and indicate that augmenting a secondary capability by a small amount can result in up to a 12% increase in robustness.

Table 9 **Top 10 players who improve robustness (RVF$^+$) through a modest increase in a capability at the secondary position indicated. (See Table 11 for correspondence between team names and abbreviations.)**

Player	Team	Position	RVF$^+$ (%)
Kyle Seager	SEA	2B	12.07
Yonder Alonso	SD	1B	9.57
Yonder Alonso	SD	3B	8.12
Alex Presley	PIT	RF	7.36
Ryan Ludwick	CIN	1B	6.84
Brett Wallace	HOU	3B	5.54
Jeff Francoeur	KC	LF	5.48
Mat Gamel	MIL	LF	5.26
Taylor Green	MIL	SS	5.17
Taylor Green	MIL	2B	5.12

At the top of the list is Kyle Seager from Seattle, who also shows up in Table 7. When nature decides to destroy capacity on the Seattle roster, she injures Dustin Ackley, who plays 2B. In response, Seager, who plays 3B, satisfies the unfulfilled demand at 2B. Augmenting Seager's capability at 2B brings his capability at 2B, 3B, and SS to approximately the same level, improving his ability to chain in the infield. Because the next best player at 2B after Seager represents a big drop in capability, adding capability to Seager at 2B results in a significant improvement in Seattle's robustness to worst-case injury. Yonder Alonso, who shows up twice in Table 9, helps mitigate the drop in production that results when nature injures San Diego's first or third baseman. He is able to do this because he becomes the second best option at 1B and 3B after the respective capability augmentations. A similar finding is true for Taylor Green, who also shows up twice in Table 9 and is used primarily to complete Milwaukee's infield chain (cf. Table 3).

A final thing to note from these results is that the vast majority of player-position augmentations did not impact worst-case performance at all. The 45 player-position augmentations that improved team robustness came from only 21 teams. That means 30% of the teams saw no impact. And among the teams that were impacted positively, only two player-position augmentations mattered on average. Without our framework to help pinpoint the meaningful capability augmentations, a manager might invest significant energy in training players at positions that will ultimately have no impact on robustness to worst-case performance.

4.5. Evaluating free agents

Free agents are players without a contract and thus can sign a new contract (typically during the off-season) with any team. Top free agents can command significant salaries, due to the limited talent pool that is available in a given off-season. Thus, careful evaluation of free agents and their potential contribution to the acquiring team is an important process that every team undertakes annually. In this section, we demonstrate the generalizability of our flexibility framework to the task of evaluating free agents. We quantify the potential gain to each of the 30 MLB teams of acquiring 10 top-ranked free agents (position players only) available during the 2012 off-season.[2] Table 10 lists the players we considered, along with the team they eventually joined.

First, we removed the free agents from the teams they signed with to simulate the situation of teams having a clean slate with which to evaluate these potential additions. Next, we considered adding each free agent in isolation to each of the 30 MLB teams in two different computational experiments. The first experiment measures the improvement in expected performance by adding a free agent, similar to the experiments from Section 4.3. For this experiment, the capability matrix $\mathbf{V}^1$ represents the base roster (minus the added free agent from Table 10, if applicable) and $\mathbf{V}^2$ represents the team with the added free agent. Using the same capability matrices as the first

[2] http://www.mlbtraderumors.com/2011/10/2012-top-50-free-agents-1.html

Table 10 Ten top-ranked free agents from the 2012 off-season. (See Table 11 for correspondence between team names and abbreviations.)

Player	2012 Team	Avg (RAR)	SD (RAR)	CoV (%)
Albert Pujols	LAA	41.7	3.3	7.8
Aramis Ramirez	MIL	17.2	3.5	20.4
Carlos Beltran	STL	18.2	3.1	16.9
Carlos Pena	TB	8.3	4.0	48.1
Coco Crisp	OAK	7.1	3.2	45.5
David Ortiz	BOS	20.5	2.2	10.9
Jimmy Rollins	PHI	17.6	5.1	29.0
Jose Reyes	MIA	31.2	5.2	16.5
Michael Cuddyer	COL	20.3	2.5	12.5
Prince Fielder	DET	24.7	4.7	19.0

Avg = average increase in team RAR due to adding player to all 30 teams; SD = standard deviation; CoV = coefficient of variation. Since David Ortiz plays DH almost exclusively, his value is averaged over American League teams only. He would not be considered as a plausible addition to National League teams, which do not have the DH position.

experiment, the second experiment measures the improvement in worst-case performance, similar to the experiments from Section 4.4.

Table 10 summarizes the average impact of adding the free agents to each team. A star player such as Albert Pujols, who has the lowest coefficient of variation, appears to generate fairly consistent value for each team. He plays several positions with skill and typically becomes the starting 1B or 3B on every team he joins. On the other hand, the added value of players like Carlos Pena and Coco Crisp varies quite a bit depending on which team they join, which is evident from their coefficients of variation being greater than 40%. Crisp, for example, is a flexible outfielder with moderate capabilities. Thus, his added value depends a lot on the quality of the existing outfielders on the team he joins.

To explore this idea more thoroughly, we contrast the value of Aramis Ramirez and Michael Cuddyer, who would be seen as fairly close substitutes. Although they exhibit both infield and outfield flexibility with comparable capabilities, they are primarily infielders. For their nominal team assignments, Ramirez spends all of his time at 3B while Cuddyer splits time between 2B and 3B. On average, Cuddyer appears to be the better player (20.3 RAR vs. 17.2 RAR). However, Figure 6 illustrates the variability in value of both players to each team. For some teams, Cuddyer is worth over 70% more than Ramirez. On the other hand, points below the diagonal correspond to teams that would generate more value from Ramirez. In fact, what is even more interesting is that if the two teams that eventually signed Ramirez (MIL) and Cuddyer (COL) had swapped them, they had the potential to both come out ahead. For Colorado, our results show that Cuddyer was worth 17.8 RAR, while Ramirez was worth 19.0 RAR. On the other hand, for Milwaukee, Ramirez was worth 15.1 RAR, while Cuddyer was worth 17.3 RAR. We further investigated how the playing times would change to understand why this swap appears to be a win-win for both teams. Consider the hypothetical situation where Milwaukee swaps Ramirez for Cuddyer. (See Milwaukee's player v_{ij} values and default playing times in Table 3.) Losing Ramirez results in Mat Gamel and Taylor Green playing 3B at a modest decrease in capability. Then, Alex Gonzalez increases his playing time to cover SS at a slight additional loss. However, Cuddyer is a much stronger alternative at 1B than Corey Hart, so there is an increase in value to having him play 1B instead. That frees Hart to spend more time at RF, resulting in a further increase over the current starter at that position, Caleb Gindl. The net impact of this chaining is over 2 RAR.

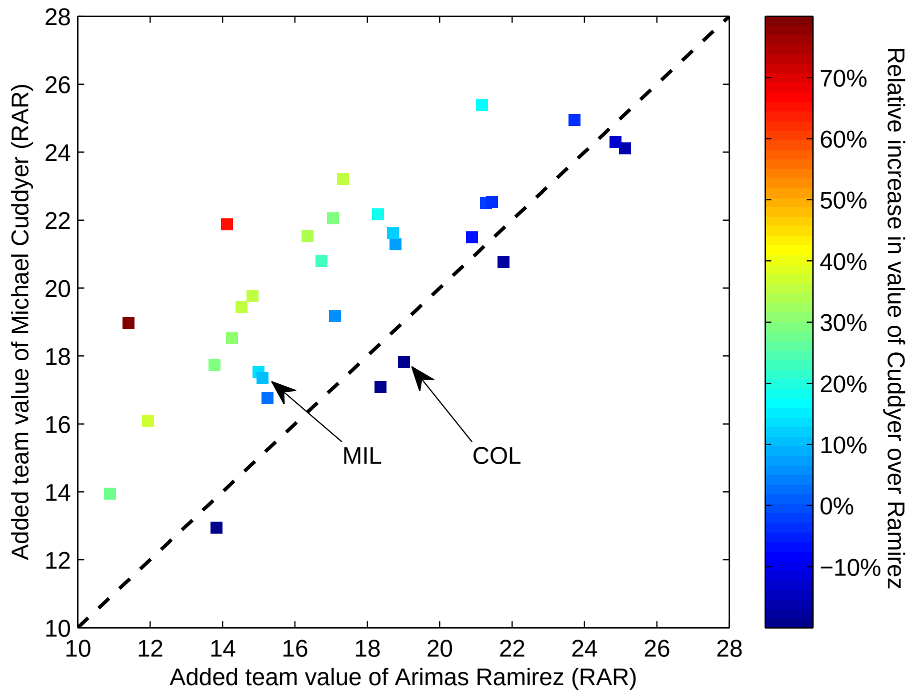

Figure 6 **Comparing the added team value (RAR) of Aramis Ramirez and Michael Cuddyer to each team.**

Finally, we briefly discuss how a robust optimization perspective can be layered on top of the analysis conducted above when evaluating free agents. Different teams have fragilities at different positions. So even though two players may have similar expected value to a team, the impact on worst-case performance can be quite different. To explore this potential effect, we quantified the change in worst-case performance of each team for each of the potential free agent additions. For a concrete example, we return to Colorado, who signed Michael Cuddyer as a free agent. It turns out three of the other players on the free agent list, if added to Colorado instead of Cuddyer, would have generated comparable average team performance (within $\sim$1 RAR): Aramis Ramirez, Carlos

Beltran, and Prince Fielder. However, each player has a unique impact on worst-case performance of the team. The change in worst-case performance is computed using the value of Γ equal to $\mathrm{PL}_{0.1}(\mathbf{V})$ for the base roster assuming Cuddyer is the free agent added ($\mathbf{V}$ is the base capability matrix). That is, we determined the smallest budget parameter in the uncertainty set required to reduce team performance by 10% in the worst-case, based on the roster where Cuddyer was the free agent addition. Then, we applied this same parameter to the new rosters with the three players substituted in for Cuddyer, one at a time, to compute three new worst-case performance levels. Figure 7 shows the change in expected performance of Colorado (x-axis) against the change in worst-case performance (y-axis) due to adding each of the three players instead of Cuddyer. As discussed above, Ramirez generates an improvement in average performance of about 1.2 RAR. He also positively affects worst-case performance by about 1.0 RAR. At the other end of the spectrum, Beltran reduces both average and worst-case team performance relative to Cuddyer. Fielder exhibits a trade-off – he has a modest positive impact on expected performance, but degrades worst-case performance by over 2 RAR. This situation is an example of a player who may look comparable or slightly better in expectation, but have hidden disadvantages that can be uncovered through our robust optimization perspective.

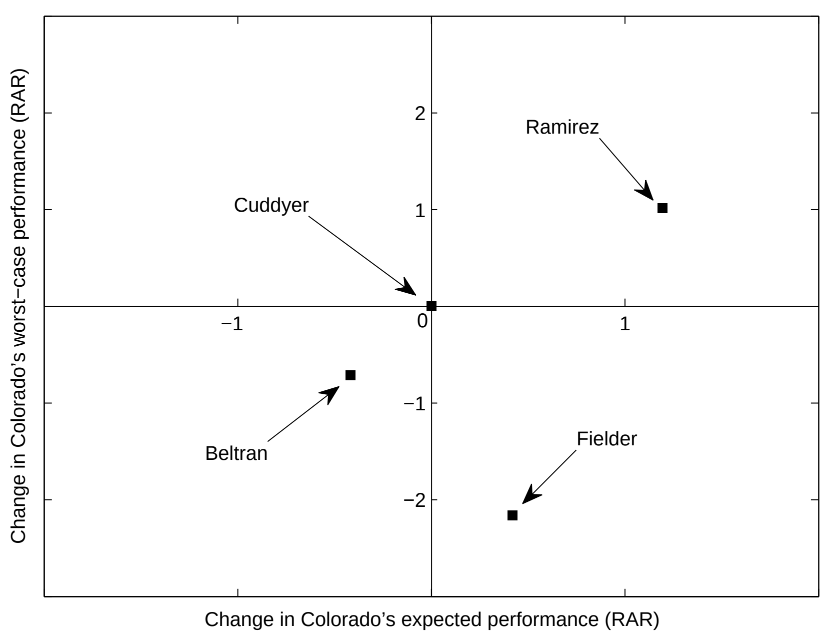

Figure 7 Change in worst-case performance of Colorado from adding players with comparable impact on expected performance.

Overall, our results illustrate how our flexibility framework can be extended to the evaluation of free agents in a comprehensive manner, considering the specific composition of the team the player would be joining. Such an analysis can be useful in distinguishing between several free agent candidates in terms of their impact on both expected and worst-case team performance. One concrete situation where such quantitative insight into player value might be important is in contract negotiations. For example, teams might identify players who would be less valuable to other teams and potentially less coveted, which could result in signing a player at a discount or having extra time to make decisions. On the flip side, a player who can demonstrate that his true value to a team is greater than other free agents on the market may have additional leverage in contract negotiations.

4.6. Who is the Most Valuable Player?

In this final case study, we present another example of the broader applicability of our flexibility framework, in this case to the quantification of "Most Valuable Players." MVP awards are ubiquitous in team sports. However, a common criticism of such an award is that it is typically given to the best performing player, not necessarily the player deemed most valuable to his or her team. Part of this mismatch is due to the fact that it is generally much easier to examine sport-specific performance metrics and determine who had superior statistical performance in a given season. However, an MVP award should really answer the question: "which player, if removed from his team, would result in the worst degradation in team performance?" We replicate our AVF experiments from Section 4.3 but with a small difference in the choice of $\mathbf{V}^1$. Instead of zeroing out a player's secondary capabilities, we zero out all capabilities, essentially substituting that player with a replacement-level player.

Figure 8 compares a team's performance loss due to losing its MVP against overall team performance. Each team's MVP is defined as the player whose removal results in the greatest degradation in team performance; we refer to the resulting degradation in performance the *MVP value*. In absolute terms, Evan Longoria would be the MVP of the league since Tampa Bay would lose the most of any team ($\sim$57 RAR) without him. We also measured relative loss, defined as the absolute loss divided by the team's overall RAR (i.e., equivalent to AVF$^-$). According to the relative metric, Ryan Zimmerman, who is worth over 24% of Washington's total production, would be the MVP. Note that one argument against using a relative measure is that it may inappropriately recognize strong players on weak teams. Thus, a third MVP metric that could be considered would be the loss (either in absolute or relative terms) conditional on the team meeting a certain level of performance such as being in the top 10 teams in the league or making the playoffs. Under this conditional performance metric, Longoria would seem to be the clear choice for MVP, as he is ranked first in both absolute and relative terms in this subset of teams.

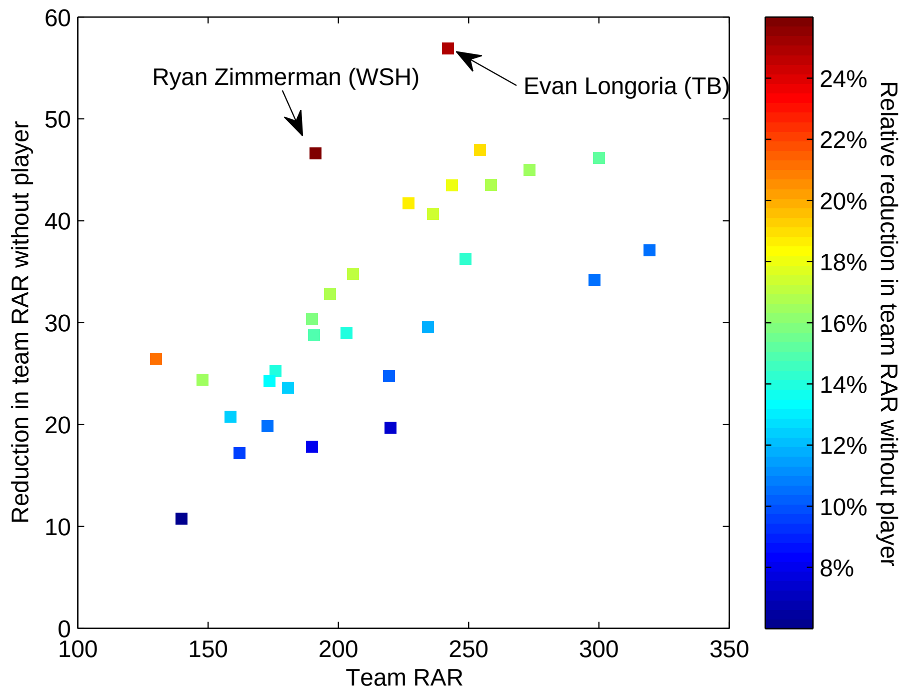

Figure 8 The value in runs above replacement of each team's MVP.

Returning to the ambiguity between "most valuable" versus "best" player, we directly compare these two quantities in Figure 9. We take the MVPs, as defined above, from each of the 30 teams and plot their MVP value (y-axis) against their individual RAR based on expected playing time over the season (x-axis). For reference, the latter quantity is what is added up over all players on a team to determine overall team performance. As expected, we observe that all points lie below the diagonal. The interpretation of this result is that a player's expected contributions to his team while playing is an upper bound on what the team loses when the player is not. If a team has capable back-up capacity, then a team will not lose anywhere near the full amount of value that a player contributes when he is playing. Points close to the diagonal in Figure 9 represent teams that do not have enough flexibility and/or capability to make up for the loss when their MVP is lost – these teams lose almost the full value that this player brings to his team. On the other hand, the points far from the diagonal represent teams with MVPs whose absence may not cripple his team as much as expected. For example, Oakland is able to recover over 50% of Yoenis Cespedes' value due to the flexibility of its roster. This finding is consistent with the results from Section 4.2, which shows that Oakland is the most robust team to worst-case performance from losing a single player. In particular, note that Cespedes is a key member of the outfield chain illustrated in Table 6, which involves flexible starters and bench players with minimal drop-off in capability. Another such example is Troy Tulowitski from Colorado. He is one of the best players in the league according to his individual RAR (∼59), but his team only loses about 45 runs (his MVP value) without him. Granted, 45 RAR is still a very large amount. However, this case is a clear example where a star player's observable statistical output, measured by his RAR in our experiments, might be a non-trivial overestimation of his value to his team. In contrast, Evan Longoria not only has high MVP value, but also high individual RAR. Thus, his high level of output when he is playing appears reflective of his true value to Tampa Bay.

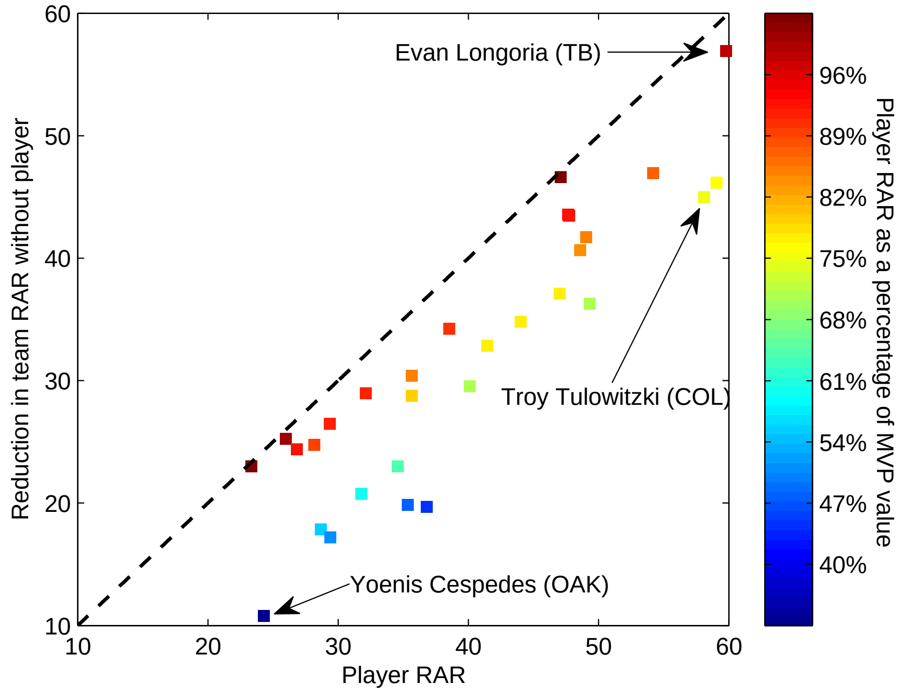

Figure 9 **Comparing MVP value against expected value in runs above replacement.**

5. Managerial Insights and Implications

In this paper, we make the first connection between process flexibility and sports analytics in the management science literature. We tackle a long-standing open question in baseball analytics, namely quantifying the value of positional flexibility, which is the ability of players to play multiple defensive positions. Our insights come from viewing a baseball team as a production network in which players are plants and positions are products. In this section, we synthesize the key findings from our computational studies that would be of interest to baseball managers, players and fans.

One of our main results is the quantification of the value of positional flexibility in terms of the average number of runs saved over the course of a season, due to flexible (and better) defenders getting more innings over bench players (Section 4.1). In short, teams may be able to attribute at least a couple wins a season to flexibility alone. On the one hand, given that average teams win about 80 games a year, a couple wins may seem modest. However, in almost every season a couple wins is the difference between making and not making the playoffs. For example, in every season since 2010, the best team to miss the playoffs was no more than two wins behind the worst playoff-bound team in its league. This finding puts into context the value that flexibility may bring to a team when searching for players to add, for example. Looking more closely at Figure 5 in Section 4.3, we see that some players' individual flexibility may be worth at least half a win (recall that 10 runs roughly corresponds to a win). Thus, one managerially useful insight is that although a team would likely never trade raw performance of a star player for an above average player with flexibility (the former would typically be worth more wins even if he played only one position, assuming he played it very well), a team on the bubble of contention may be able to strategically piece together just enough fractions of wins to make the playoffs with flexible players sourced through the trade or free agent market.

We showed how our framework can be used to evaluate the fit between players and teams with respect to flexibility and injury. For example, if two players are comparable in offensive ability, we can facilitate the complex evaluation of how a team can respond to injuries differently with these two players. Such an approach is useful when evaluating free agents, trades, or even minor league players. For example, through our free agent analysis in Section 4.5, we showed that players with comparable average value added can have quite different impact on worst-case team performance. Overall, player valuations can vary substantially across teams. In fact, our framework identified a case where it appeared that if two teams swapped their major free agent signings, they could have both come out ahead.

By adding a robust optimization lens through which to view flexibility, our approach can also be used to identify teams that are susceptible to worst-case injuries (Section 4.2). To address such vulnerabilities, our approach can be used prescriptively to indicate which players to add to optimize worst-case performance. Or, as shown in Section 4.4, our robust model can identify players on an existing team to train in a secondary position to minimize the impact of worst-case injury. For example, we showed that augmenting the capability of a specific player at a specific position can improve team robustness to worst-case injury by over 12%. For teams with limited trade or free agent options, internal improvement may be the most cost-effective option. With access to minor league rosters, a manager could extend our approach to identify minor leaguers to develop or players to draft. This type of succession planning would be especially useful for teams with star players who have a history of injury, which will only get worse with age.

Our framework enables us to quantify exactly how much value individual players can add to their team via their flexibility. For example, we find that flexibility alone of the "Most Flexible

Player" (Section 4.3) is single-handedly is worth 2.5% of his team's production, which represents almost one third of the total value of flexibility for his team. We also verify that while individual flexibility always benefits the team, for almost a quarter of the players evaluated, their flexibility actually hurts their own (defensive) performance. Finally, we showed how our flexibility-measuring framework can be extended to address the general determination of "Most Valuable Player" in a rigorous, data-driven manner (Section 4.6). Objectively determining the MVP is difficult for two reasons. First, while an MVP award is meant to identify the player who is most valuable to his team, it instead is typically awarded to someone who simply had the best offensive statistics, because the latter is easier to appreciate. Second, even if the intent is to measure MVP appropriately, it is hard to quantify how the team would have performed without the player in question. This type of "with or without you" analysis is increasingly common in sports analytics, but prior to this paper such a rigorous determination of the performance of a team with or without a player, considering replacement level alternatives and how players would shift defensive positions accordingly, was missing. By extending the value of flexibility measurement to measuring the value of having the player on the roster at all, we could determine what a team loses in performance without a specific player. Furthermore, we showed that this loss can be much less than the player's expected contributions, which reinforces the idea that simply looking at statistical output does not always provide an accurate characterization of a player's true value to his team. Finally, novel metrics like our MVP metric can also increase fan engagement, as fans and writers are constantly looking for new insights into player performance that can be debated.

One potential implication of our work is the development of new player valuation metrics that could impact contracts. We envision such metrics being useful for players and managers when it comes to (re-)negotiating contracts, as they can be used by stakeholders on both sides to strengthen their case in a novel way. For example, our flexibility results could help a player make the case that he contributes more value to his team than his individual statistics might suggest (being in the upper left of Figure 5). Flexibility value is harder to quantify that just looking at offensive statistics and it might be something a player gets less credit for. Most contract negotiations leverage simple comparables analysis with similar players using simple statistics; our approach would add another level of depth to such negotiations.

Appendix A: Baseball background

In baseball, the objective is to win each game by scoring more *runs* than the opposing team. Play proceeds in *innings* with each team taking a turn on offense (*at bat*) and defense (*in the field*). The team that scores the most runs after nine innings is the winner – if the score is tied, extra innings are added to the game one at a time until the tie is broken.

Prior to each game, teams choose which players will start the game (the *starting lineup*), order these players for batting (the *batting order*), and assign these players to positions in the field. See Figure 10 for a depiction of the playing field and the nine fielding positions: pitcher (P), catcher (C), first baseman (1B), second baseman (2B), third baseman (3B), shortstop (SS), left fielder (LF), center fielder (CF), and right fielder (RF).

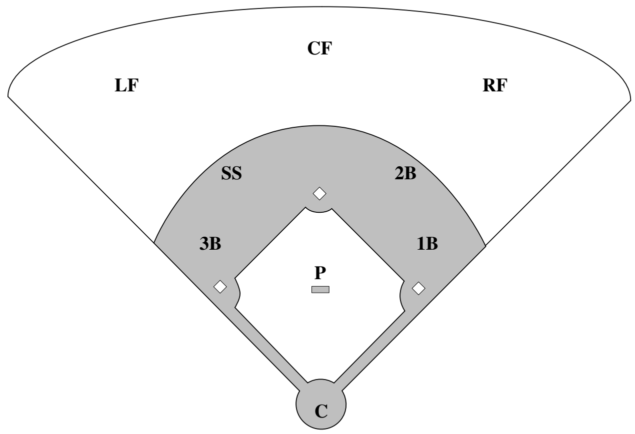

Figure 10 Fielding positions in a baseball field.

When a team is at bat, players bat in the specified order. The team's turn at bat ends when three of the batting team's players have been declared out by the *umpires*. In the following inning, the team re-starts at the spot in the batting order where it previously left off. A run is scored when a player makes it back to home plate by traversing the bases in a counter-clockwise fashion without being tagged or forced out. Thus, each player contributes to the team's production on offense through effective batting (i.e., reaching and traversing the bases) and on defense through effective fielding (i.e., creating outs and limiting movement around the bases). The one exception is due to the *designated hitter rule*, which allows teams to designate another player to take the place of the pitcher in the batting order. Thus, the designated hitter does not contribute to defense.[3]

Each team maintains a 25-man roster of players who are available to play on a daily basis. Of these players, approximately 11 to 12 are pitchers. In addition to the 25-man roster, each team maintains an extended 40-man roster that includes the 25-man roster plus 15 additional players who can be readily added to the 25-man roster, for example in the case of injuries. Players who are injured may be added to the *disabled list* (DL) to free up roster spots. The regular season in baseball comprises 162 games. Top teams from the regular season make the playoffs (the *postseason*) and vie for the championship (the *World Series*[4]). We refer the reader to the publication *Official Baseball Rules* (Major League Baseball 2015) for a comprehensive description of the game.

Table 11 lists the current teams in Major League Baseball, along with their abbreviations and logos.

[3] The designated hitter rule was adopted by the American League, one of the two sub-leagues in Major League Baseball, in 1973. All games played in an American League team's stadium apply the designated hitter rule.

[4] The championship series in Major League Baseball is referred to as the World Series, although 29 of the 30 teams are located in the United States and the remaining is located in Canada.

Table 11 Major League Baseball teams

League	Division	Name	Abbrev.	Logo
American	Eastern	Baltimore Orioles	BAL	
		Boston Red Sox	BOS	
		New York Yankees	NYY	
		Tampa Bay Rays	TB	
		Toronto Blue Jays	TOR	
	Central	Chicago White Sox	CWS	
		Cleveland Indians	CLE	
		Detroit Tigers	DET	
		Kansas City Royals	KC	
		Minnesota Twins	MIN	
	Western	Houston Astros	HOU	
		Los Angeles Angels	LAA	
		Oakland Athletics	OAK	
		Seattle Mariners	SEA	
		Texas Rangers	TEX	
National	Eastern	Atlanta Braves	ATL	
		Miami Marlins	MIA	
		New York Mets	NYM	
		Philadelphia Phillies	PHI	
		Washington Nationals	WSH	
	Central	Chicago Cubs	CHC	
		Cincinnati Reds	CIN	
		Milwaukee Brewers	MIL	
		Pittsburgh Pirates	PIT	
		St. Louis Cardinals	STL	
	Western	Arizona Diamondbacks	ARI	
		Colorado Rockies	COL	
		Los Angeles Dodgers	LAD	
		San Diego Padres	SD	
		San Francisco Giants	SF	

Appendix B: Model estimation

In this section we detail the data gathered and models developed for estimating player capacities (Section B.1) and capabilities (Section B.2). This section is largely taken from the conference paper Chan and Fearing (2013).

B.1. Capacity model

To determine capacity statistics, we compared player-days on the DL to player-days on the major league roster from 1999 (the earliest available DL data) to 2011 (the last year before our study) using data obtained from MLB's eBIS database, the electronic Baseball Information System available to all 30 MLB teams. Note that this comparison should result in an optimistic estimate of player capacities, because short-term injuries often do not result in a stint on the DL. To estimate the distribution of DL days, we considered the impact of player age and the primary position played through parameterized statistical modeling. In the various models we tested, the age of the player consistently ranked as a statistically significant factor, whereas the player's primary position did not. In particular, age was a highly significant factor in determining the probability that a player would enter the DL in a given season. We subsequently measure the duration the player spent on the DL as a percentage of the days spent on the major league roster. One interesting finding is that conditioned on entering the DL, the total duration is positively-skewed, and not significantly impacted by age. This observation leads us to the two-stage statistical model described in equations (8) and (9), where a DL occurrence is Bernoulli distributed and its duration is log-normally distributed. In equation (8), $p_{jt} \in \{0,1\}$ represents the probability player j joins the DL in season t and x_{jt}^{age} equal the player's age in days on July 1 of season t. In equation (9), $y_{jt} \in (0,1]$ represents the proportion of the season spent on the DL and $\epsilon_{jt} \sim N(0,\sigma^2)$. In our training set, we exclude players who spent less than 80 days on the major league roster in a season to protect against selection biases (e.g., young players sent back to the minors after experiencing an injury).

$$\ln\left(\frac{p_{jt}}{1-p_{jt}}\right) = \beta_0 + \beta^{\text{age}} x_{jt}^{\text{age}}, \tag{8}$$

$$\ln(y_{jt}) = \beta_0 + \epsilon_{jt}, \tag{9}$$

Thus, the first model estimates the probability of going on the DL at least once in a season, and the second estimates the total duration of DL stints as a percentage of the season. We estimate these models using the R functions glm and lm, respectively. The corresponding parameter estimates and statistical tests are provided below, along with a plot comparing the two-stage model estimates to the training data.

Table 12 DL occurrence (equation (8))

	Estimate	Std. Error	z value	$\Pr(>\lvert z\rvert)$
Intercept	-1.195e+00	1.414e-01	-8.452	$<$ 2e-16
AgeDays	6.885e-05	1.298e-05	5.305	1.13e-07

Table 13 DL duration (equation (9))

	Estimate	Std. Error	z value	$\Pr(>\lvert z\rvert)$
Intercept	-1.25255	0.01271	-98.51	$<$ 2e-16

In Figure 11, the lines plot the two-stage model estimates (at the mean (black), 90th (green), and 95th (red) percentiles) and the circles plot the empirical averages for the player-seasons within the training data matching the corresponding age bucket (i.e., age ± 0.5 years). Our two-stage model appears to be slightly more conservative regarding the tails of the injury distribution, compared to the mean of the distribution, which is estimated quite accurately.

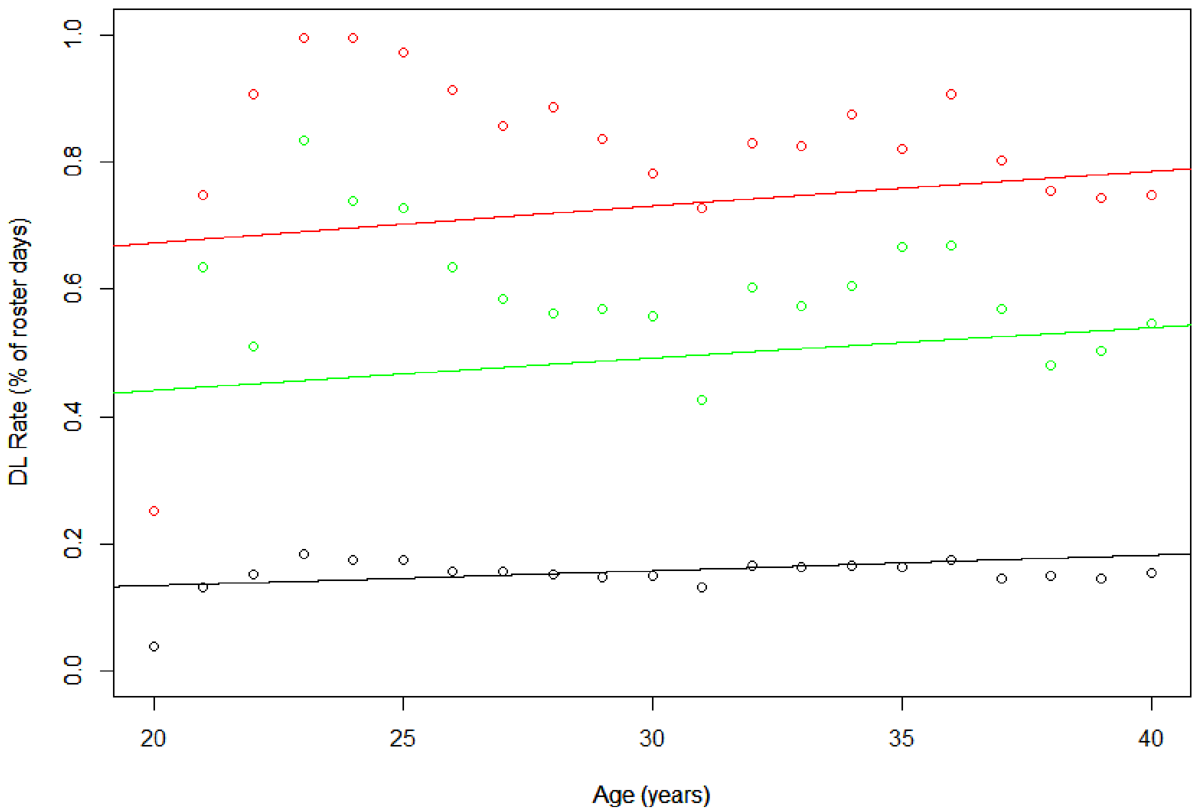

Figure 11 The relationship between DL rate and player age.

B.2. Capability model

The capability of a player at a position (v_{ij}) is composed of three parts: 1) the player's expected offensive contribution in runs above average per game, 2) a position-specific, replacement-level adjustment per game, and 3) the player's expected defensive contribution in runs above average per game. For the offensive contributions, we collect projections of offensive statistics for each player using the ZiPS projection system (Szymborski 2012), which we then convert into runs per game using linear weights (FanGraphs 2013b). The conversion is based on estimating the value of particular events such as a walk, hit, home run, etc. We used ZiPS projections rather than any other projection system because the ZiPS projections covered a larger number of the roster players considered in our study. For positional adjustments, we rely on the values used by FanGraphs (FanGraphs 2013c). These adjustments enable comparisons in production between players of different positions and are necessary since the offensive and defensive contributions are calculated in terms of runs above average for that specific position. To adjust these values for in-season replacements, which we assume are less readily available, we set the replacement level at each position (i.e., v_{ij} equal to zero) lower by an additional 10 runs per season. Thus, for example, instead of applying a position adjustment of +2.5 runs per season for a second baseman (FanGraphs 2013c), we apply an adjustment of +12.5 runs per season.

For defensive capabilities, we use the UZR/150 statistic (FanGraphs 2013d), where UZR stands for "Ultimate Zone Rating". The Ultimate Zone Rating quantifies the number of runs saved by a player's defensive ability and UZR/150 is a normalized value roughly corresponding to the number of opportunities expected in 150 games of playing time. The statistic is not calculated for catchers, so for the purposes of our study, we assume that all catchers have a UZR/150 equal to zero. We gathered this data for the 2012 season from FanGraphs (FanGraphs 2013a).

The most challenging data requirement for our model is the need to provide expected defensive contributions for each player at each position he could play, even if he has never actually done so. Defensive statistics are known to be quite noisy, often requiring multiple years of data to stabilize. In order to address this instability, especially for player-position combinations with few innings played, we assume that a player's defensive capabilities across all positions are drawn from a multivariate normal population distribution. That is, we assume that 1) there is a normal distribution of ability at each position across all major league players, and 2) there is a correlation structure that describes the relationship between abilities at one position and any other. The model is described in equations (10), (11), and (12). The dependent variable y_{ijt} represents the UZR/150 for player j at position i in season t. We weight each observation by n_{ijt}, the number of innings played at the position. We classify positions into one of three groups – infield, outfield, and catcher/first base – which form the set G. In addition to the player's age (x_{jt}^{age}), this model incorporates the following features:

- x_{jt}^{base}: factor variable indicating the base position for the player – the position most frequently played during the prior four seasons, or the ZiPS position for players with no prior major league experience,
- x_{ijt}^{pos}: factor variable equal to the position i for which the UZR/150 is being estimated,
- $x_{ijt}^{\text{left}\times\text{pos}}$: factor variable equal to the position for players that throw left-handed, and a default value for right-handed throwers,
- $x_{ijt}^{g_{\text{base}}\times g_{\text{pos}}}$: factor variable indicating the fielding group transition being made; positions are categorized into three groups: infield, outfield, or catcher/first-base,
- x_{ijt}^{min}: binary variable indicating if player i has played position j at least nine innings during the prior four seasons,
- x_{ijt}^{exp}: variable indicating the natural logarithm of the innings in excess of 90 played at the position during the prior four seasons; equal to 0 if fewer than 90 innings have been played.

Given these variables, the model is:

$$\begin{aligned} y_{ijt} = {} & \beta^{\text{base}} x_{jt}^{\text{base}} + \beta^{\text{pos}} x_{ijt}^{\text{pos}} + \beta^{g_{\text{base}}\times g_{\text{pos}}} x_{ijt}^{g_{\text{base}}\times g_{\text{pos}}} + \beta^{\text{left}\times\text{pos}} x_{ijt}^{\text{left}\times\text{pos}} \\ & + \sum_{g\in G} (\beta_g^{\text{age}} x_{jt}^{\text{age}} + \beta_g^{\text{min}} x_{ijt}^{\text{min}} + \beta_g^{\text{exp}} x_{ijt}^{\text{exp}}) + \alpha_{ij} + \epsilon_{ijt}, \end{aligned} \tag{10}$$

$$\epsilon_{ijt} \sim N\left(0, \frac{\sigma_j^2}{n_{ijt}}\right), \tag{11}$$

$$\alpha_{ij} \sim N(0, \mathbf{\Sigma}). \tag{12}$$

Note that the variance in the error term ϵ_{ijt} is inversely proportional to the number of innings played at the position whereas the population distribution has a covariance matrix $\mathbf{\Sigma}$ that is independent of the innings played. Thus, in addition to providing the correlations we seek, this approach implicitly regularizes the data; that is, the more innings a player has played, the further his UZR/150 estimates can deviate from the population mean. In order to address selection biases that otherwise skew the results, we subtract five runs (approximately half a win) from the observed UZR/150 values for players in the training data who have no prior experience at a position ($x_{ijt}^{\min} = 0$). We estimate this model using lmer, an R package for incorporating random effects into generalized linear regression, and the parameter estimates are provided in Tables 14 and 15.

Table 14 Fixed effects coefficients

	Estimate	Std. Error	t value
BasePosition1B+C	-5.486e+00	2.053e-01	-26.72
BasePosition2B	-7.095e-01	2.641e-01	-2.69
BasePosition3B	-2.714e-01	2.608e-01	-1.04
BasePositionCF	5.443e+00	2.442e-01	22.28
BasePositionLF	-5.516e-01	2.415e-01	-2.28
BasePositionRF	1.241e-01	2.425e-01	0.51
BasePositionSS	2.128e+00	2.649e-01	8.03
Position2B	-7.437e+00	3.003e-01	-24.77
Position3B	-8.107e+00	2.968e-01	-27.32
PositionCF	1.154e+01	3.029e-01	38.08
PositionLF	1.859e+01	3.036e-01	61.24
PositionRF	1.851e+01	3.040e-01	60.91
PositionSS	-1.011e+01	3.010e-01	-33.59
GroupTransitionIF-1B	-6.455e+00	1.829e-01	-35.30
GroupTransitionIF-OF	2.592e+00	2.230e-01	11.62
GroupTransitionOF-1B	-5.486e+00	1.569e-01	-34.97
GroupTransitionOF-IF	-3.475e+00	2.443e-01	-14.22
ThrowsPositionLeftL-1B	1.420e+00	6.413e-02	22.15
ThrowsPositionLeftL-CF	-1.425e+00	8.825e-02	-16.14
ThrowsPositionLeftL-LF	2.338e+00	8.987e-02	26.01
ThrowsPositionLeftL-RF	-1.104e+00	9.814e-02	-11.25
PositionGroup1B:HasExperience	6.349e+00	1.134e-01	55.98
PositionGroupIF:HasExperience	6.626e+00	8.509e-02	77.88
PositionGroupOF:HasExperience	8.319e+00	7.455e-02	111.60
PositionGroup1B:ExcessExperience	2.042e-02	9.007e-03	2.27
PositionGroupIF:ExcessExperience	4.320e-01	5.841e-03	73.97
PositionGroupOF:ExcessExperience	4.124e-02	6.479e-03	6.36
PositionGroup1B:AgeDays	-1.060e-04	1.708e-05	-6.21
PositionGroupIF:AgeDays	-1.848e-04	1.324e-05	-13.96
PositionGroupOF:AgeDays	-2.499e-03	1.637e-05	-152.66

Table 16 lists the top 10 players at each position excluding catcher and their corresponding UZR/150 estimates based on our model. Our model suggests that strong defensive players will be defensively capable at multiple positions. For example, Nick Punto, a shortstop, also shows up as a top 10 second and third baseman. Similarly, many of the top 10 center fielders also place in the top 10 in both left and right field.

Table 15 Random effects covariance matrix

	1B	2B	3B	CF	LF	RF	SS
1B	0.05645	0.04139	0.08828	0.07053	0.00236	0.05221	0.08168
2B	0.04139	0.08392	0.10947	0.07049	0.04122	0.13485	0.06527
3B	0.08828	0.10947	0.20168	0.06445	-0.03412	0.10595	0.12832
CF	0.07053	0.07049	0.06445	0.23878	0.18258	0.23105	0.11307
LF	0.00236	0.04122	-0.03412	0.18258	0.30420	0.20186	0.00011
RF	0.05221	0.13485	0.10595	0.23105	0.20186	0.34971	0.09855
SS	0.08168	0.06527	0.12832	0.11307	0.00011	0.09855	0.12305

Table 16 Top 10 players in UZR/150 by position

Rank	1B	2B	3B	SS	LF	CF	RF
1	M. Kotsay	C. Utley	N. Punto	N. Punto	B. Gardner	F. Gutierrez	B. Carroll
2	C. Kotchman	N. Punto	E. Longoria	A. Everett	T. Campana	T. Gwynn	T. Gwynn
3	A. Pujols	P. Polanco	A. Beltre	J.J. Hardy	N. Morgan	P. Bourjos	F. Gutierrez
4	I. Davis	J. McDonald	C. Counsell	B. Ryan	T. Gwynn	C. Gomez	A. Torres
5	K. Youkilis	B. Zobrist	S. Rolen	C. Izturis	A. Torres	B. Carroll	T. Campana
6	D. Barton	D. Pedroia	P. Polanco	E. Andrus	C. Gomez	N. Morgan	N. Morgan
7	C. Utley	C. Counsell	A. Everett	J. McDonald	S. Cousins	T. Campana	C. Gomez
8	A. Jones	B. Phillips	J. Hannahan	C. Barmes	J. Ellsbury	A. Torres	B. Zobrist
9	M. Teixeira	M. Ellis	R. Zimmerman	P. Janish	B. Carroll	B. Gardner	B. Revere
10	J. Morneau	J. Uribe	J. McDonald	C. Counsell	R. Johnson	B. Revere	S. Cousins

Appendix C: An integer programming reformulation of the robust player-position assignment model

Dualizing the inner maximization problem, formulation (4) becomes the following bilinear optimization model:

$$
\begin{aligned}
\underset{\mathbf{c},\mathbf{p},\mathbf{q}}{\text{minimize}} \quad & \sum_{i\in I} d_i p_i + \sum_{j\in J} c_j q_j \\
\text{subject to} \quad & p_i + q_j \geq v_{ij},\ i\in I, j\in J, \\
& \sum_{j\in J} \frac{u_j - c_j}{\sigma_j} = \Gamma, \\
& l_j \leq c_j \leq u_j,\ j\in J, \\
& p_i \geq 0,\ i\in I, \\
& q_j \geq 0,\ j\in J.
\end{aligned}
\tag{13}
$$

While global optimization techniques may be used to solve formulation (13), we derive an equivalent mixed-integer linear programming formulation that can be solved using a standard MIP solver. Note that without loss of generality, we may bound $p_i + q_j$ from above by $v_{\max} = \max_{i,j} v_{ij}$. As a bilinear problem in the variables $\mathbf{c}$ and $(\mathbf{p},\mathbf{q})$, the corresponding bounded polyhedral feasible regions $\mathbf{C}_\Gamma$ and $\mathbf{P} = \{(\mathbf{p},\mathbf{q}) \mid v_{\max} \geq p_i + q_j \geq v_{ij},\ i\in I, j\in J\}$ are disjoint. It is well-known that in this case, there exists an optimal solution $(\mathbf{c}^*,\mathbf{p}^*,\mathbf{q}^*)$ to formulation (13) such that $\mathbf{c}^*$ is a vertex of $\mathbf{C}_\Gamma$ and $(\mathbf{p}^*,\mathbf{q}^*)$ is a vertex of $\mathbf{P}$ (Horst et al. 2000).

To derive an equivalent MIP for formulation (13) and introduce the appropriate binary variables, we need to characterize the structure of the vertices of $\mathbf{C}_\Gamma$. Notice that $\mathbf{C}_\Gamma$ is simply the intersection of a hyperrectangle with a single equality constraint. Thus, a vertex $\mathbf{c} \in \mathbf{C}_\Gamma$ will have $c_j = l_j$ for $j \in J_1$, $c_j = u_j$ for $j \in J_2$ and $c_j = (\sum_{k \in J_1}(u_k - l_k)/\sigma_k - \Gamma)\sigma_j + u_j$ if $j \in J_3$ (a singleton set), where J_1, J_2, J_3 form a partition of J. We can now reformulate (13) into a mixed-integer program using binary variables to decide which indices are chosen to form J_1 and J_3. The following formulation introduces binary variables (with appropriate auxiliary variables and constraints) to optimize over the vertices of $\mathbf{C}_\Gamma$.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{i \in I} d_i p_i + \sum_{j \in J} u_j q_j - \sum_{j \in J}(u_j - l_j) Q_j^z - \Gamma \sum_{j \in J} \sigma_j Q_j^w + \sum_{j \in J} \sum_{k \in J} \frac{\sigma_j (u_k - l_k)}{\sigma_k} Q_{jk}^{wz} \\
\text{subject to} \quad & p_i + q_j \geq v_{ij},\ i \in I, j \in J, \\
& \left(\Gamma w_j - \sum_{k \in J} \frac{z_k (u_k - l_k)}{\sigma_k} \right) \leq \frac{u_j - l_j}{\sigma_j},\ j \in J, \\
& \sum_{j \in J} \frac{z_j (u_j - l_j)}{\sigma_j} \leq \Gamma, \\
& w_j \leq (1 - z_j),\ j \in J, \\
& \sum_{j \in J} w_j = 1, \\
& Q_j^z \leq q_j,\ j \in J, \\
& Q_j^z \leq z_j v_{\max},\ j \in J, \\
& Q_j^w \leq q_j,\ j \in J, \\
& Q_j^w \leq w_j v_{\max},\ j \in J, \\
& Q_{jk}^{wz} \geq q_j - v_{\max}(2 - z_k - w_j),\ j \in J, k \in J, \\
& Q_{jk}^{wz} \leq w_j v_{\max},\ j \in J, k \in J, \\
& Q_{jk}^{wz} \leq z_k v_{\max},\ j \in J, k \in J, \\
& p_i \geq 0,\ i \in I, \\
& q_j, Q_j^z, Q_j^w \geq 0,\ j \in J, \\
& Q_{jk}^{wz} \geq 0,\ j \in J, k \in J, \\
& z_j, w_j \in \{0, 1\},\ j \in J.
\end{aligned}
\tag{14}
$$

Lemma 1. *Formulation* (13) *is equivalent to formulation* (14).

Proof of Lemma 1 Let $z_j = 1$ if $c_j = l_j$ and 0 otherwise. Let $w_j = 1$ if $c_j = (\sum_{k \in J_1}(u_k - l_k)/\sigma_k - \Gamma)\sigma_j + u_j$ and 0 otherwise. We reformulate the constraint $\sum_{j \in J}(u_j - c_j)/\sigma_j = \Gamma$ as

$$\sum_{j \in J} \frac{z_j (u_j - l_j) + \delta_j}{\sigma_j} = \Gamma$$

and add the constraints $\sum_{j \in J} w_j = 1$, $w_j \leq 1 - z_j$ for $j \in J$, and $\delta_j \leq w_j (u_j - l_j)$ for $j \in J$. Together, these constraints ensure that δ_j is nonzero for only one j (when $w_j = 1$) and takes the appropriate value to ensure the equality constraint involving Γ is satisfied. We may now replace c_j with $u_j - z_j(u_j - l_j) - \delta_j$ and derive an equivalent formulation to (4) where the outer minimization is over the vertices of $\mathbf{C}_\Gamma$:

$$\begin{aligned}
\min_{\mathbf{z},\mathbf{w},\boldsymbol{\delta}} \max_{\mathbf{x}} \quad & \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij} \\
\text{subject to} \quad & \sum_{j\in J} x_{ij} \le d_i,\ i\in I, \\
& \sum_{i\in I} x_{ij} \le u_j - z_j(u_j - l_j) - \delta_j,\ j\in J, \\
& \delta_j \le w_j(u_j - l_j),\ j\in J, \\
& w_j \le (1 - z_j),\ j\in J, \\
& \sum_{j\in J} w_j = 1, \\
& \sum_{j\in J} \frac{z_j(u_j - l_j) + \delta_j}{\sigma_j} = \Gamma, \\
& x_{ij} \ge 0,\ i\in I, j\in J, \\
& \delta_j \ge 0,\ j\in J, \\
& z_j, w_j \in \{0,1\},\ j\in J.
\end{aligned} \tag{15}$$

The equality constraint involving Γ can be written as

$$\delta_j = \sigma_j \left(\Gamma - \sum_{k\in J} \frac{z_k(u_k - l_k)}{\sigma_k} - \sum_{k\neq j} \frac{\delta_k}{\sigma_k}\right), \tag{16}$$

which is equivalent to

$$\delta_j = w_j\sigma_j \left(\Gamma - \sum_{k\in J} \frac{z_k(u_k - l_k)}{\sigma_k} - \sum_{k\neq j} \frac{\delta_k}{\sigma_k}\right), \tag{17}$$

since w_j is binary and $w_j = 0$ forces $\delta_j = 0$ by the third constraint. Note that when $w_j = 1$, $w_k = 0$ for all $k \neq j$ and thus $\delta_k = 0$ for all $k \neq j$. Thus, constraint (17) is equivalent to

$$\delta_j = w_j\sigma_j \left(\Gamma - \sum_{k\in J} \frac{z_k(u_k - l_k)}{\sigma_k}\right). \tag{18}$$

Non-negativity of δ_j is easily re-formulated based on (18). Lastly, the constraint $\delta_j \le w_j(u_j - l_j)$ can be re-written as

$$w_j \left(\Gamma - \sum_{k\in J} \frac{z_k(u_k - l_k)}{\sigma_k}\right) \le w_j \left(\frac{u_j - l_j}{\sigma_j}\right),\ j\in J, \tag{19}$$

or

$$\left(\Gamma w_j - \sum_{k\in J} \frac{z_k(u_k - l_k)}{\sigma_k}\right) \le \frac{u_j - l_j}{\sigma_j},\ j\in J, \tag{20}$$

since w_j is binary.

Thus, formulation (15) can be re-written without the δ_j variables:

$$
\begin{aligned}
\min_{\mathbf{z},\mathbf{w}} \max_{\mathbf{x}} \quad & \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij} \\
\text{subject to} \quad & \sum_{j\in J} x_{ij} \le d_i,\ i\in I, \\
& \sum_{i\in I} x_{ij} \le u_j - z_j(u_j-l_j) - w_j\sigma_j\left(\Gamma - \sum_{k\in J}\frac{z_k(u_k-l_k)}{\sigma_k}\right),\ j\in J, \\
& \left(\Gamma w_j - \sum_{k\in J}\frac{z_k(u_k-l_k)}{\sigma_k}\right) \le \frac{u_j-l_j}{\sigma_j},\ j\in J, \\
& \sum_{j\in J}\frac{z_j(u_j-l_j)}{\sigma_j} \le \Gamma, \\
& w_j \le (1-z_j),\ j\in J, \\
& \sum_{j\in J} w_j = 1, \\
& x_{ij}\ge 0,\ i\in I, j\in J, \\
& z_j, w_j \in\{0,1\},\ j\in J.
\end{aligned}
\tag{21}
$$

Dualizing the inner maximization in formulation (21), we get the following mixed-integer bilinear problem

$$
\begin{aligned}
\underset{\mathbf{z},\mathbf{w},\mathbf{p},\mathbf{q}}{\text{minimize}} \quad & \sum_{i\in I} d_i p_i + \sum_{j\in J}\left(u_j - z_j(u_j-l_j) - w_j\sigma_j\left(\Gamma - \sum_{k\in J}\frac{z_k(u_k-l_k)}{\sigma_k}\right)\right) q_j \\
\text{subject to} \quad & p_i + q_j \ge v_{ij},\ i\in I, j\in J, \\
& w_j\left(\Gamma - \sum_{k\in J}\frac{z_k(u_k-l_k)}{\sigma_k}\right) \le w_j\left(\frac{u_j-l_j}{\sigma_j}\right),\ j\in J, \\
& \sum_{j\in J}\frac{z_j(u_j-l_j)}{\sigma_j} \le \Gamma, \\
& w_j \le (1-z_j),\ j\in J, \\
& \sum_{j\in J} w_j = 1, \\
& p_i \ge 0,\ i\in I, \\
& q_j \ge 0,\ j\in J, \\
& z_j, w_j\in\{0,1\},\ j\in J.
\end{aligned}
\tag{22}
$$

Finally, we may reformulate (22) into a mixed-integer *linear* problem. Recall we may bound $p_i + q_j$ (and because of non-negativity, both p_i and q_j individually) above by $v_{\max} = \max_{i,j} v_{ij}$. Using the constraints $Q_j^z \le q_j$ and $Q_j^z \le z_j v_{\max}$, we enforce $Q_j^z = z_j q_j$ at an optimal solution. Using the constraints $Q_j^w \le q_j$ and $Q_j^w \le w_j v_{\max}$, we enforce $Q_j^w = w_j q_j$ at an optimal solution. Using the constraints $Q_{jk}^{wz} \ge q_j - v_{\max}(2 - z_k - w_j)$, $Q_{jk}^{wz} \le w_j v_{\max}$ and $Q_{jk}^{wz} \le z_k v_{\max}$, we enforce $Q_{jk}^{wz} = w_j z_k q_j$ at an optimal solution. The direction of the first inequality in each set is due to the sign of the corresponding term in the objective function. Adding these new variables and constraints into the model, we arrive at the final formulation (14). □

C.1. A special case and a greedy solution

Under certain assumptions, we can analytically characterize an optimal solution to the outer minimization in formulation (4). This characterization allows us to identify exactly how nature will injure players – that is, in what order will players be injured and how much will each of their capacities be reduced. We consider the special case where the capability values are independent of the players and dependent only on the position (i.e., v_{ij} depends only on i).

PROPOSITION 1. *Without loss of generality, assume $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n$. Let $v_{ij} = v_i$ for all i,j and $\Gamma \in [0, \sum_{j\in J}(u_j - l_j)/\sigma_j]$. Let k satisfy $\sum_{j=1}^{k-1}(u_j - l_j)/\sigma_j < \Gamma \leq \sum_{j=1}^{k}(u_j - l_j)/\sigma_j$. An optimal solution $\mathbf{c}^*$ to the outer minimization in formulation* (4) *satisfies*

$$c_j^* = \begin{cases} l_j, & j = 1, \dots k-1, \\ u_j - \sigma_j \left(\Gamma - \sum_{j=1}^{k-1}(u_j - l_j)/\sigma_j\right), & j = k, \\ u_j, & j = k+1, \dots, n. \end{cases} \tag{23}$$

Given nature's optimal decision $\mathbf{c}^*$ from Proposition 1, we can also characterize the optimal response $\mathbf{x}^*$.

PROPOSITION 2. *Let $v_{ij} = v_i$ for all i,j, and without loss of generality assume $v_1 \geq v_2 \geq \cdots \geq v_m$. Let $\mathbf{c}^*$ be given by* (23). *If $\sum_{j\in J} c_j^* > \sum_{i\in I} d_i$, an optimal solution $\mathbf{x}^*$ to the inner minimization problem in formulation* (4) *is $x_{ij}^* = d_i c_j^* / \sum_{j\in J} c_j^*$ for all i,j. If $\sum_{j\in J} c_j^* \leq \sum_{i\in I} d_i$, and t satisfies $\sum_{i=1}^{t-1} d_i < \sum_{j\in J} c_j^* \leq \sum_{i=1}^{t} d_i$, then for $j = 1, \dots, n$,*

$$x_{ij}^* = \begin{cases} d_i c_j^* / \sum_{j\in J} c_j^*, & i = 1, \dots t-1, \\ \left(\sum_{j\in J} c_j^* - \sum_{i=1}^{t-1} d_i\right) c_j^* / \sum_{j\in J} c_j^*, & i = t, \\ 0, & i = t+1, \dots, m. \end{cases} \tag{24}$$

These results state that when capabilities are dependent only on the position being played, nature will greedily injure players in order of decreasing σ_j, (reducing c_j to l_j, before moving on to c_{j+1}) in order to maximize utilization of the budget of uncertainty Γ. In response, the team greedily assigns players to positions in descending order of capability value, respecting whichever is the limiting constraint, total capacity or total demand. While the assumption that capabilities are only position-dependent is somewhat unrealistic and a simplification, this results provides insight into how a team might greedily prioritize the allocation of playing time.

Proof of Proposition 1 Note that at an optimal solution to formulation (13), $q^* = \max_i\{v_i - p_i^*, 0\}$. Substituting in this expression, we get

$$\begin{aligned} \underset{\mathbf{c}}{\text{minimize}} \quad & \sum_{i\in I} d_i p_i^* + \max_i\{v_i - p_i^*, 0\} \sum_{j\in J} c_j \\ \text{subject to} \quad & \sum_{j\in J} \frac{u_j - c_j}{\sigma_j} = \Gamma, \\ & l_j \leq c_j \leq u_j,\ j \in J, \end{aligned} \tag{25}$$

which is equivalent to a continuous knapsack problem with variables $(u_j - c_j)$. Thus, the equation for $\mathbf{c}^*$ follows immediately. $\square$

Proof of Proposition 2 With $\mathbf{c}^*$ fixed, the inner maximization is

$$\begin{aligned} \underset{\mathbf{x}}{\text{maximize}} \quad & \sum_{i\in I} v_i \sum_{j\in J} x_{ij} \\ \text{subject to} \quad & \sum_{j\in J} x_{ij} \leq d_i,\ i \in I, \\ & \sum_{i\in I} x_{ij} \leq c_j^*,\ j \in J, \\ & x_{ij} \geq 0,\ i \in I, j \in J. \end{aligned} \tag{26}$$

If we sum over j in the second constraint and replace $\sum_{j \in J} x_{ij}$ with y_i throughout, we get

$$\begin{aligned} \underset{\mathbf{y}}{\text{maximize}} \quad & \sum_{i \in I} v_i y_i \\ \text{subject to} \quad & y_i \leq d_i,\ i \in I, \\ & \sum_{i \in I} y_i \leq \sum_{j \in J} c_j^*, \\ & y_i \geq 0,\ i \in I, \end{aligned} \tag{27}$$

from which it becomes obvious that the greedy solution is optimal. It is straightforward to check that the expressions for x_{ij}^* satisfy the constraints in both cases. $\square$

Acknowledgments

We are grateful to the department editor, associate editor and reviewers for comments and guidance that greatly improved the paper. The authors thank Justin Cho for helping to collect and process the MLB performance data, and Justin Boutilier for helping with some of the computational experiments.

References

Alderson, D. L., G. G. Brown, W. M. Carlyle, R. K. Wood. 2011. Solving defender-attacker-defender models for infrastructure defense. R. K. Wood, R. F. Dell, eds., *Operations Research, Computing, and Homeland Defense*. INFORMS, Hanover, MD, 28–49.

Bandi, C., D. Bertsimas. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Math. Prog.* **134**(1) 23–70.

Bassamboo, A., R. S. Randhawa, J. A. Van Mieghem. 2011. A little flexibility is all you need: On the asymptotic value of flexible capacity in parallel queuing systems. *Oper. Res.* **60**(6) 1423–1435.

Brown, G., M. Carlyle, J. Salmeron, K. Wood. 2006. Defending critical infrastructure. *Interfaces* **36**(6) 530–544.

Brown, G. G., W. M. Carlyle, R. C. Harney, E. M. Skroch, R. K. Wood. 2009. Interdicting a nuclear-weapons project. *Oper. Res.* **57**(4) 866–877.

Brusco, M., T. Johns. 1998. Staffing a multiskilled workforce with varying levels of productivity: An analysis of cross-training policies. *Decision Sciences* **29**(2) 499–515.

Chan, T. C. Y., D. S. Fearing. 2013. The value of flexibility in baseball roster construction. *Proceedings of the 2013 MIT Sloan Sports Analytics Conference*.

Chou, M., C.-P. Teo, H. Zheng. 2008. Process flexibility: Design, evaluation, and applications. *Flexible Services and Manufacturing J.* **20**(1) 13–58.

Chou, M., C.-P. Teo, H. Zheng. 2011. Process flexibility revisited: The graph expander and its applications. *Oper. Res.* **59**(5) 1090–1105.

Chou, M. C., G. A. Chua, C.-P. Teo. 2010a. On range and response: Dimensions of process flexibility. *Eur. J. Oper. Res.* **207**(2) 711–724.

Chou, M. C., G. A. Chua, C-P. Teo, H. Zheng. 2010b. Design for process flexibility: Efficiency of the long chain and sparse structure. *Oper. Res.* **58**(1) 43–58.

Deng, T., Z.-J. M. Shen. 2013. Process flexibility design in unbalanced networks. *Manufacturing Service Oper. Management* **15**(1) 24–32.

FanGraphs. 2013a. 2012 UZR. URL `http://www.fangraphs.com/leaders.aspx?pos=all&stats=fld&lg=all&qual=0&type=1&season=2012&month=0&season1=2012&ind=0&team=0&rost=0&age=0&filter=&players=0,accessedJanuary10,2013`.

FanGraphs. 2013b. Linear Weights. URL `http://www.fangraphs.com/library/index.php/principles/linear-weights/,accessedJanuary10,2013`.

FanGraphs. 2013c. Positional Adjustment. URL `http://www.fangraphs.com/library/misc/war/positional-adjustment/,accessedJanuary10,2013`.

FanGraphs. 2013d. Ultimate Zone Rating (UZR). URL `http://www.fangraphs.com/library/defense/uzr/,accessedJanuary10,2013`.

FanGraphs. 2014. What is WAR? URL `http://www.fangraphs.com/library/misc/war/`.

Ganeshapillai, G., J. Guttag. 2014. A data-driven method for in-game decision making in MLB. *Proceedings of the 2014 MIT Sloan Sports Analytics Conference*.

Graves, S. C., B. T. Tomlin. 2003. Process flexibility in supply chains. *Management Sci.* **49**(7) 907–919.

Gurumurthi, S., S. Benjaafar. 2004. Modeling and analysis of flexible queuing systems. *Naval Res. Logist.* **51**(5) 755–782.

Hopp, W. J., S. M. R. Iravani, W. L. Xu. 2010. Vertical flexibility in supply chains. *Management Sci.* **56**(3) 495–502.

Hopp, W. J., E. Tekin, M. P. Van Oyen. 2004. Benefits of skill chaining in serial production lines with cross-trained workers. *Management Sci.* **50**(1) 83–98.

Horst, R., P. M. Pardalos, N. V. Thoai. 2000. *Introduction to Global Optimization*. 2nd ed. Kluwer.

Iravani, S. M., M. P. Van Oyen, K. T. Sims. 2005. Structural flexibility: A new perspective on the design of manufacturing and service operations. *Management Sci.* **51**(2) 151–166.

James, B. 1977. *Baseball Abstract*. Self-published.

Jordan, W. C., S. C. Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management Sci.* **41**(4) 577–594.

Major League Baseball. 2015. Official Baseball Rules. URL `http://mlb.mlb.com/mlb/downloads/y2015/official_baseball_rules.pdf`.

Mak, H.-Y., Z.-J. M. Shen. 2009. Stochastic programming approach to process flexibility design. *Flexible Services and Manufacturing J.* **21**(3) 75–91.

Simchi-Levi, D., H. Wang, Y. Wei. 2015. Increasing supply chain robustness through process flexibility and inventory. *Under review* .

Simchi-Levi, D., Y. Wei. 2012. Understanding the performance of the long chain and sparse designs in process flexibility. *Oper. Res.* **60**(5) 1125–1141.

Simchi-Levi, D., Y. Wei. 2015. Worst-case analysis of process flexibility designs. *Oper. Res.* **63**(1) 166–185.

Streib, N., S. J. Young, J. Sokol. 2011. A Major League Baseball team uses Operations Research to improve draft preparation. *Interfaces* **42** 119–130.

Szymborski, D. 2012. 2012 ZiPS projections spreadsheets, v.1. URL `www.baseballthinkfactory.org/oracle/discussion/2012_zips_projections_spreadsheets_v._1`.

Tang, C., B. Tomlin. 2008. The power of flexibility for mitigating supply chain risks. *Int. J. Production Economics* **116**(1) 12–27.

Trick, M. A., H. Yildiz, T. Yunes. 2011. Scheduling Major League Baseball umpries and the traveling umpire problem. *Interfaces* **42** 232–244.

Tsitsiklis, J. N., K. Xu. 2012. On the power of (even a little) resource pooling. *Stochastic Systems* **2**(1) 1–66.

Wallace, R. B., W. Whitt. 2005. A staffing algorithm for call centers with skill-based routing. *Manufacturing Service Oper. Management* **7**(4) 276–294.

Wang, X., J. Zhang. 2015. Process flexibility: A distribution-free bound on the performance of k-chain. *Oper. Res.* **63** 555–571.

Woolner, K. 2000. Baseball's Hilbert problems. C. Kahrl, C. Davenport, J. S. Sheehan, R. Jazayerli, eds., *Baseball Prospectus 2000*. Potomac Books Inc.

Wyers, C. 2010. Manufactured Runs. URL `http://www.baseballprospectus.com/article.php?articleid=10917`.

Yan, Z., S. Y. Gao, C. P. Teo. 2017. On the design of sparse but efficient structures in operations. *Management Sci.* **In Press**.

Zimmerman, J. 2013. 2013 Disabled List Team Data. URL `http://www.fangraphs.com/blogs/2013-disabled-list-team-data/,accessedMay9,2016`.