The value of flexibility in baseball roster construction

Timothy C. Y. Chan
University of Toronto
Toronto, ON, Canada, M5S 3G8
Email: tcychan@mie.utoronto.ca

Douglas S. Fearing
Harvard Business School
Cambridge, MA, USA, 02134
Email: dfearing@hbs.edu

Abstract

Drawing inspiration from the theory of production flexibility in manufacturing networks, we provide the first optimization-based analysis of the value of positional flexibility (the ability of a player to play multiple positions) for a major league baseball team in the presence of injury risk. First, we develop novel statistical models to estimate (1) the likelihood and duration of player injuries during the regular season, and (2) fielding abilities at secondary fielding positions. Next, we develop a robust optimization model to calculate the degradation in team performance due to injuries. Finally, we apply this model to measure the difference in performance between a team with players who have positional flexibility and a team that does not. We find that using 2012 rosters, flexibility was expected to create from 3% (White Sox) to 15% (Cubs) in value for each team, measured in runs above replacement. In analyzing the results, we find that platoon advantages (e.g., having left-handed batters face right-handed pitchers) form an important component of flexibility. As a secondary finding, based on our statistical analysis of injuries, we find that the likelihood of injury increases with age, but the duration of injury does not.

1 Introduction

The most important decisions facing a sports team often relate to roster construction and playing time allocation. Such problems are complex and, in the context of trades or free agent acquisitions, must account for factors such as the existing team composition, salary cap/luxury tax implications, and short-term versus long-term team performance. In baseball, the roster construction problem is further complicated by the fact that players can typically play multiple defensive positions; such positional swaps occur regularly, especially towards the end of close games when teams introduce pinch hitters/runners. While positional flexibility–a player's ability to play multiple positions–introduces additional dynamics into the roster construction problem, it can also contribute to a team's competitive advantage. A baseball team lacking positional flexibility may be reluctant to bring in a pinch hitter in the top half of the inning if there is not enough flexibility to field a decent defense in the bottom half of the inning.

Although it provides a team with options, flexibility alone may not be enough. A team needs flexibility and capability–that is, players with the ability to play multiple positions <u>with skill</u>. Furthermore, the value of flexibility is not restricted to player substitutions in a game. Teams with flexible players are also better equipped to deal with player injuries that occur randomly throughout the season. Thus, flexibility is also related to a team's depth, in terms of having skilled bench players who can fill in as needed.

The manner of flexibility we focus on in this paper has been studied quite extensively in the manufacturing industry [1]. Consider a car manufacturer with multiple factories and multiple product lines. Suppose each factory is capable of producing only one type of car and each car type can be produced in only one factory. If there is a sudden loss of production capacity in a factory (e.g., machine breakdown), the company may be unable to produce enough cars to meet demand. On the other hand, if the factories are flexible enough to produce multiple car types and each car type can be produced at multiple factories, then the company would be better able to respond to changes in supply or demand. Baseball players are like factories producing innings-played at each position.

In this paper, we borrow ideas from manufacturing flexibility and apply it to baseball roster construction. We will use optimization models to evaluate the flexibility of Major League Baseball teams in a rigorous way and quantify the value of flexibility to each team using publicly available performance data. Our specific contributions are: (1) we develop new analytical methods to estimate injury rates and the effect of changing the fielding position a player takes, (2) we develop a new metric based on a robust optimization model to assess roster vulnerabilities, and (3) we quantify the value of flexibility for Major League Baseball teams in 2012.

2 Basic player-position assignment model

First, we formulate a model that determines the assignment of players to positions in order to maximize the performance of the team. Suppose player j plays position i with a "capability" value of v_{ij}. The player's capability represents the amount the player contributes to the team's performance (e.g., runs) per game by playing that position. A capability value of 0 represents a within-season, readily-available, replacement-level player. Player j can play up to c_j innings per season, and each position i must be assigned a player for up to d_i innings per season (all unassigned innings are assumed at replacement level). Letting x_{ij} represent the fraction of all innings in a season that player j is assigned to position i, assume that a team has J players and there are I positions to be filled. In this case, we can find the optimal assignment of players to positions by solving the following linear optimization model:

$$
\begin{aligned}
\text{Max}_x \quad & \sum_{i=1}^{I} \sum_{j=1}^{J} v_{ij} x_{ij} \\
\text{Subject to} \quad & \sum_{i=1}^{I} x_{ij} \leq c_j, \ \text{ for all } j = 1, \dots, J \\
& \sum_{j=1}^{J} x_{ij} \leq d_i, \ \text{ for all } i = 1, \dots, I \\
& x_{ij} \geq 0, \ \text{ for all } i = 1, \dots, I, j = 1, \dots, J
\end{aligned}
\tag{1}
$$

As mentioned above, this is a standard assignment model that is commonly used in manufacturing to optimize the production of goods, with factories instead of players and products instead of positions. To further tailor this model to baseball, we adjust the "demand" constraints to account for the handedness of the opposing pitcher. That is, the number of innings that position i needs to be assigned a player, d_i, can be split into two quantities d_i^R and d_i^L to account for right- and left-handedness of the opposing pitcher, which results in two (structurally identical) demand constraints instead of one. Additionally, to limit the number of innings a player can be assigned to play catcher, we add a cost of $(1/0.85)$ to these assignments in the capacity constraint, which ensures that catchers are assigned to play at most 85% of the innings in a season. This model has several implicit assumptions. First, it assumes that capability values are additive across players. This is a fairly common assumption in baseball research, but would be less appropriate in other sports. Second, injuries (which determine the player "capacity" values c_j) are

realized before the assignment of player-innings to positions.   Last, the timing in which players are injured during the season is controlled by the team. This last assumption is much stronger but is necessary for model tractability.

3   Robust player-position assignment model

Not all injuries are equally devastating to a team. Imagine if "nature" could consciously choose which players to injure and for what duration of time in order to maximally degrade a team's performance. By conducting such an experiment, we would be able to evaluate the value that flexibility provides to a team–a more flexible team will be better-equipped to respond, even if nature is consciously deciding which players to injure. Thus, to measure the value of positional flexibility, we augment model (1) to allow an extra decision (by nature) to determine the worst-case combination of player injuries. This is our robust optimization model, which is shown as model (2). For ease of exposition, we leave out the baseball-specific extensions.

$$
\begin{aligned}
&\text{Min}_c \text{ Max}_x && \sum_{i=1}^{I} \sum_{j=1}^{J} v_{ij} x_{ij} \\
&\text{Subject to} && \sum_{i=1}^{I} x_{ij} \le c_j, \ \text{ for all } j = 1, \dots, J \\
& && \sum_{j=1}^{J} x_{ij} \le d_i, \ \text{ for all } i = 1, \dots, I \\
& && \sum_{j=1}^{J} \frac{u_j - c_j}{\sigma_j} \le \Gamma \\
& && l_j \le c_j \le u_j, \ \text{ for all } j = 1, \dots, J \\
& && x_{ij} \ge 0, \ \text{ for all } i = 1, \dots, I, j = 1, \dots, J
\end{aligned}
\tag{2}
$$

In the robust model, nature is allowed to choose the player capacities first (the outer minimization over the decision variables c_j) and then the player-position assignments are made afterwards (the inner maximization over the decision variables x_{ij}). Recall that in model (1), the fixed capacity parameter c_j incorporates an estimate of how many innings player j will miss due to injury. The robust model considers a distribution of c_j values (with standard deviation σ_j, upper bound u_j and lower bound l_j) corresponding to a range of injury scenarios. Nature is allowed to injure players as long as the resulting c_j values stay within the bounds defined by the third and fourth set of constraints in model (2)–an "uncertainty set" [2]. The value Γ represents the "budget of disruption" that nature is given. It is used to couple injuries across players and limit the total amount of injuries a team can face–the larger the Γ, the more injuries nature can introduce. We solve model (2) by transforming it into an equivalent mixed-integer linear optimization problem and using the commercial solver CPLEX.

Model (2) allows us to evaluate the worst-case performance of a team in the presence of injury risk. The result is worst-case because nature is allowed to decide which players to injure to minimize team performance. Another useful quantity is average-case team performance, which we can measure by simulating injuries. Specifically, we randomly generate many c_j values from an injury distribution (whose construction is described in Section 4) and then solve the resulting model (1) with those c_j values.

To perform our computations, many of the parameters in the model need to be specified as inputs. These parameters include the team rosters, which determine the scope of each model; the player capacity statistics (σ_j, u_j, l_j); and the player-position capabilities (v_{ij}). In the next section, we describe how these data are collected and processed.

4   Data

To determine which players to include in each team's optimization models, we used the 2012 opening day rosters, including players on the 25-man roster, 40-man roster, and major league disabled list (DL) at the beginning of the season. The data was collected from [3]. Some teams may incorporate additional flexibility by using non-roster players signed to minor league contracts, but this is outside the scope of our study.

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). Note that this 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 leads us to the two-stage statistical model described in Equations (3) and (4), where a DL occurrence is Bernoulli distributed and its duration is log-normally distributed. In Equation (3), we let

$p_{jt} \in \{0,1\}$ indicate 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 (4), we let $y_{jt} \in (0,1]$ indicate the proportion of the season spent on the DL. 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).

$$\ln\left(\frac{p_{jt}}{1-p_{jt}}\right) = \beta_0 + \beta_{age}x_{jt}^{\text{age}} \qquad (3)$$

$$\ln(y_{jt}) = \beta_0 + \epsilon_{jt}, \quad \epsilon_{jt} \sim \mathcal{N}(0, \sigma^2) \qquad (4)$$

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 in the Appendix, along with a plot comparing the two-stage model estimates to the training data.

Within the assignment 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 use Dan Szymborski's ZiPs projections for each team [4], which we convert to runs per game using linear weights [5]. 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. The fact that ZiPs projections are broken down by opposing handedness and outcome also makes them easy to convert into the required runs per game metrics. For positional adjustments we rely on the values used by FanGraphs and developed by Tom Tango [6]. To adjust these 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, we apply an adjustment of +12.5 runs per season.

For defensive capabilities, we start with the UZR/150 statistic provided by FanGraphs [7]. The statistic is intended to normalize defensive contributions by position to roughly correspond 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. 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 (5). For the dependent variable, we let y_{ijt} represent the UZR/150 for player j at position i in season t, and we weight each observation by n_{ijt}, the number of innings played at the position. In addition to the player's 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 4 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} : (0, 1) variable indicating if player i has played position j at least 9 innings during the prior 4 seasons,
- x_{ijt}^{exp} : variable indicating the natural-logarithm of the innings in excess of 90 played at the position during the prior 4 seasons, and equal to 0 if fewer than 90 innings have been played.

$$y_{ijt} = \beta_{x_{jt}^{\text{base}}}^{\text{base}} + \beta_{x_{ijt}^{\text{pos}}}^{\text{pos}} + \beta_{x_{ijt}^{g_{\text{base}} \times g_{\text{pos}}}}^{g_{\text{base}} \times g_{\text{pos}}} + \beta_{x_{ijt}^{\text{left} \times \text{pos}}}^{\text{left} \times \text{pos}} + \sum_g \left[\beta_g^{\text{age}}x_{jt}^{\text{age}} + \beta_g^{\text{min}}x_{ijt}^{\text{min}} + \beta_g^{\text{exp}}x_{ijt}^{\text{exp}}\right] + \alpha_{ij} + \epsilon_{ijt},$$

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

$$\alpha_j \sim \mathcal{N}(0, \boldsymbol{\Sigma}) \qquad (5)$$

Note that in Equations (5) the variance on the error term ϵ_{ijt} is inversely proportional to the number of innings played at the position whereas the population distribution has a variance-covariance matrix $\mathbf{\Sigma}$ that is independent of the innings played. Thus, in addition to providing the correlations we seek, this 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 5 runs (or 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 the Appendix. Table 10 in the Appendix lists the top 10 players at each position (excluding catcher) and their corresponding UZR/150 estimates based on our model.

5 Data analysis and results

In this section, we consider multiple scenarios to demonstrate the power of our approach and the metrics that it supports. In each case, we compare an assignment with the full set of capabilities (full) to one with a subset of these capabilities representing a team with no flexibility or depth (no-flex). To determine the no-flex capabilities, we first assign players to positions using Model (1) without injuries and without considering splits by opposing pitcher handedness. That is, each player is given a single capability value at each position which assumes the player faces a left-handed pitcher 29.8% of the time (the overall split from the 2012 season). This results in the assignment of at most two players to catcher and one player to each other position. To build the no-flex capabilities matrix, we zero out all capabilities corresponding to positions the player was not assigned to. For example, in the combined model, Matt Kemp is assigned to play centerfield (CF). Thus in the no-flex capabilities matrix, his v_{ij} values will be set equal to zero for all i that do not correspond to CF.

In the first scenario, we evaluate no-injury assignments with and without flexibility. This corresponds to solving Model (1) with two capability matrices (full and no-flex) for each team. In addition to the total runs above replacement (RAR), Table 1 lists the percentage gain in performance based on incorporating flexibility.

Table 1: The value of flexibility in the absence of injuries

Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)	Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)	Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)
LAD	171	182	5.8	MIN	198	203	2.4	STL	241	243	1.0
SDP	175	184	4.9	MIL	235	241	2.3	TEX	332	335	1.0
CHC	198	208	4.5	BOS	353	361	2.2	KC	211	212	0.5
PHI	270	281	3.9	WAS	211	216	2.1	BAL	217	218	0.5
NYM	187	194	3.3	TOR	281	286	1.7	TBR	280	281	0.4
PIT	187	193	3.3	DET	254	258	1.6	CIN	282	282	0.1
CLE	220	228	3.2	SEA	143	145	1.3	MIA	260	260	0.0
OAK	154	159	3.0	CWS	201	203	1.1	COL	303	303	0.0
SFG	212	218	2.8	ARI	252	255	1.1	NYY	340	340	0.0
LAA	266	273	2.7	ATL	226	229	1.0	HOU	166	166	0.0

In the no-injury case, flexibility provides value by allowing teams to exploit platoon advantages (i.e., left-handed batters facing right-handed pitchers). In Table 2 we illustrate this fact by comparing the position assignments for the Los Angeles Dodgers (LAD), the team with the highest value of flexibility in the no-injury case. We exclude the players with no assignments in either model. Full flexibility allows the Dodgers to introduce left-right platoons at 1B, 2B, and RF. Notice that in the no-flex case, each player is assigned one position and each position is assigned one player (except for catcher due to the 85% assumption stated earlier).

Table 2: Position assignment comparison for LAD with and without flexibility in the absence of injuries. The color indicates the proportion of time each player is assigned to each position against left-handed pitchers (green) and right-handed pitchers (red).

	No flexibility	RAR	With flexibility	RAR
A.J. Ellis	C (0.25/ 0.60)	17.9	C (0.25/ 0.60)	29.4
T. Federowicz	C (0.04/ 0.11)	1.4	C (0.04)	4.7
J. Loney	1B (0.30/0.70)	10.5	1B (0.70)	12.8
A. Kennedy	2B (0.30/0.70)	13.0	2B (0.70)	17.4
J. Uribe	3B (0.30/0.70)	22.9	3B (0.30/0.70)	14.6

	No flexibility	RAR	With flexibility	RAR
J. Hairston	SS (0.30/0.70)	18.5	SS (0.30/0.70)	28.9
A. Ethier	LF (0.30/0.70)	23.6	LF (0.70)	20.7
M. Kemp	CF (0.30/0.70)	38.0	CF (0.30/0.70)	10.1
T. Gwynn	RF (0.30/0.70)	25.5	RF (0.70); LF (0.30)	15.4
M. Treanor			C (0.11)	2.3
J. Rivera			1B (0.30)	9.6
M. Ellis			2B (0.30)	7.1
J. Sands			RF (0.30)	5.8

In the second scenario, we simulate injuries by sampling player capacities based on the injury models described in Section 4. Then, we solve Model (1) for both the full and no-flex cases using the sampled c_j values to determine the assignments with and without flexibility, respectively. This procedure is performed 250 times for each team and set of capabilities, and the results are averaged across these iterations. A comparison of the averaged results (full vs. no-flex) is provided in Table 3.

Table 3: The value of flexibility in the presence of simulated injuries

Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)	Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)	Team	No flex (RAR)	W/flex (RAR)	Value of flex (%)
CHC	167	197	15.0	PIT	158	175	9.9	WAS	179	191	6.4
LAD	143	164	12.8	ARI	214	237	9.9	DET	214	227	5.5
NYM	155	176	11.8	STL	199	220	9.9	NYY	283	299	5.2
OAK	130	147	11.7	CIN	236	261	9.5	SEA	124	131	5.1
PHI	223	253	11.7	LAA	224	245	8.5	KC	181	190	4.9
SDP	150	170	11.4	COL	252	273	7.9	ATL	192	202	4.8
SFG	175	196	10.8	TEX	276	299	7.5	TBR	235	247	4.8
CLE	185	207	10.7	MIA	217	234	7.5	BAL	181	189	4.4
MIL	200	224	10.4	TOR	239	257	7.3	HOU	141	146	3.5
BOS	295	327	10.0	MIN	168	180	6.7	CWS	171	177	3.4

As with the previous comparison, there is significant variation in the value of flexibility across teams, ranging from 3.4% for the Chicago White Sox (CWS) to 15.0% for the Chicago Cubs (CHC). Table 4 compares the position assignments for CHC with and without flexibility. As above, when there is no flexibility, each player is assigned one position and each position is assigned one player (except catcher). Under no flexibility, the total playing time assigned to each position is less than one, with the excess innings valued at the replacement level. As with the Dodgers, the no-injuries model suggests left-right platoons for the Cubs at 3 positions – 2B, LF, and CF. These platoons persist in the presence of simulated injuries as seen in Table 4. Additionally, the Cubs have a relatively balanced starting lineup with a deep bench, helping to mitigate the impact of an injury to any player. For example, at 3B against a right-handed pitcher, the model suggests that the Cubs could play Stewart (+0.15 runs / game) or LaHair (+0.14 runs / game). Although LaHair has never played 3B in the majors, our fielding model estimated his 3B skill at -14.5 UZR / 150, a few runs below the worst defensive third-basemen – Ty Wiggins with a -12.4 UZR / 150 and Chris Johnson with a -11.6 UZR / 150.

Table 4: Position assignment comparison for CHC with and without flexibility in the presence of simulated injuries. The color indicates the proportion of time each player is assigned to each position against left-handed pitchers (green) and right-handed pitchers (red).

	No flexibility	RAR	With flexibility	RAR
G. Soto	C (0.21/0.49)	28.8	C (0.21/0.50)	29.4
S. Clevenger	C (0.08/0.19)	6.1	C (0.01/0.18)	4.7
A. Rizzo	1B (0.24/0.57)	12.4	1B (0.25/0.59)	12.8
B. DeWitt	2B (0.25/0.58)	16.7	2B (0.04/0.58); 3B (0.21/0.02)	17.4
I. Stewart	3B (0.25/0.60)	18.6	3B (0.08/0.46); 2B (0.08); SS (0.05)	14.6
S. Castro	SS (0.26/0.61)	28.4	SS (0.26/0.60)	28.9
D. DeJesus	LF (0.24/0.56)	24.5	LF (0.01/0.57); RF (0.07/0.02)	20.7
D. Sappelt	CF (0.24/0.56)	15.2	CF (0.25/0.24)	10.1
M. Byrd	RF (0.24/0.56)	15.6	RF (0.22/0.54); CF (0.01/0.01)	15.4
W. Castillo			C (0.08/0.02)	2.3

	No flexibility	RAR	With flexibility	RAR
A. Soriano			LF (0.21/0.12)	9.6
B. LaHair			3B (0.22); 1B (0.01/0.10); RF (0.03)	7.1
J. Baker			2B (0.25); RF (0.01)	5.8
L. Valbuena			2B (0.04)	0.5
D. Barney			SS (0.04/0.04); 2B (0.01/0.01)	1.5
R. Johnson			LF (0.08); CF (0.03)	3.2
T. Campana			CF (0.44); RF (0.11); LF (0.01)	12.7

In the third scenario, we evaluate player-position assignments with and without flexibility when nature is allowed to determine the worst-case combination of player injuries (i.e., Model (2)). In the formulation of the robust assignment model, Γ represents the sum of the standard deviations of lost capacity across the team. For a single value of Γ, we can perform the same comparisons as above to determine the value of flexibility with respect to a specific level of disruption. Alternatively, we can define the team's protection level (PL) with respect a certain loss level as the minimum value of Γ that is required to force the team to incur such a loss. For example, we would define a team's 5% PL as the smallest value of Γ that results in a 5% loss in RAR from the no-injury assignment. A team that is more robust to injuries would thus have a higher PL value. In Table 5, we list the 5% and 10% PL values for each team[1].

Table 5: Robust 5% and 10% protection levels

Team	5% PL	10% PL	Team	5% PL	10% PL	Team	5% PL	10% PL
OAK	2.70	5.02	STL	1.34	2.76	COL	0.99	1.97
MIL	1.93	3.89	LAD	1.31	2.77	HOU	0.98	1.97
CHC	1.82	4.16	ATL	1.19	2.41	CIN	0.95	1.90
PIT	1.53	3.25	TEX	1.16	2.33	SFG	0.94	1.89
SDP	1.52	3.25	CLE	1.09	2.19	TOR	0.94	1.88
ARI	1.49	3.37	PHI	1.08	2.17	DET	0.86	1.72
BOS	1.43	2.92	CWS	1.04	2.09	LAA	0.83	1.66
NYY	1.43	2.97	KC	1.04	2.09	SEA	0.80	1.60
MIN	1.40	2.77	BAL	1.02	2.03	TBR	0.72	1.45
NYM	1.38	2.81	MIA	1.00	2.00	WAS	0.66	1.31

According to this metric, the Nationals' and Rays' rosters are the least robust to injuries. In each case, the lack of robustness is due to the presence of a superstar third-baseman (Ryan Zimmerman and Evan Longoria, respectively), whose injury would result in approximately replacement-level players moving into the lineup (Danny Espinosa / Adam LaRoche for the Nationals and Jeff Keppinger / Elliot Johnson for the Rays). The Athletics's roster is the most robust according to this metric. This is in part based on the balance of the lineup; Yoenis Cespedes is estimated to be the most productive player, at 29 RAR for the season. Additionally, if Cespedes is injured, the players pushed into the lineup as a result (Michael Taylor / Kila Ka'aihue) would be able to cover almost two-thirds of the resulting loss.

6 Conclusion and extensions

In this paper we developed a new and highly-tractable optimization-based approach for evaluating baseball rosters in the presence of uncertainty due to player injury risks. Using this approach, we calculated the value of flexibility in different contexts and compared these values across teams. We believe our novel identification of the variation in flexibility and resulting value represents an under-explored opportunity for teams to improve their rosters.

There are multiple directions for future work. This paper focused on measuring the flexibility of a team, but the same optimization model can be adapted to quantify the value of flexibility each player provides his team. The injury model can be augmented with player-specific injury histories. Minor league position experience can also be incorporated into the fielding model. A team interested in building flexibility may consider spring training positional experiments, the results of which can be incorporated into the optimization model to determine the added benefit of the developed flexibility.

---

[1] PL values are calculated by iteratively solving Model (2), increasing Γ by 0.25 at each step, and applying a linear interpolation of the Γ values that bound the specified loss.

7 Acknowledgments

8 References

[1] W. C. Jordan and S. C. Graves, "Principles on the benefits of manufacturing process flexibility," *Management Science*, vol. 41, no. 4, pp. 577-594, 1995.

[2] C. Bandi and D. Bertsimas, "Tractable stochastic analysis in high dimensions via robust optimization," *Mathematical Programming*, vol. 134, no. 1, pp. 23-70, 2012.

[3] www.mlbdepthcharts.com. Accessed January 10, 2013. <http://www.mlbdepthcharts.com/>.

[4] Dan Szymborski. "2012 ZiPS Projections Spreadsheets, v.1." Baseball Think Factory. Accessed January 10, 2013, <www.baseballthinkfactory.org/oracle/discussion/2012_zips_projections_spreadsheets_v._1>.

[5] Graham MacAree. "Linear Weights". FanGraphs. Accessed January 10, 2013. <http://www.fangraphs.com/library/index.php/principles/linear-weights/>.

[6] Graham MacAree. "Positional Adjustment". FanGraphs. Accessed January 10, 2013. <http://www.fangraphs.com/library/index.php/misc/war/positional-adjustment/>.

[7] www.fangraphs.com. Accessed January 10, 2013. <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>.

9 Appendix

Parameter estimates for the player capacity model

Table 6: DL occurrence coefficients (equation (3))

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

Table 7: DL duration coefficients (equation (4))

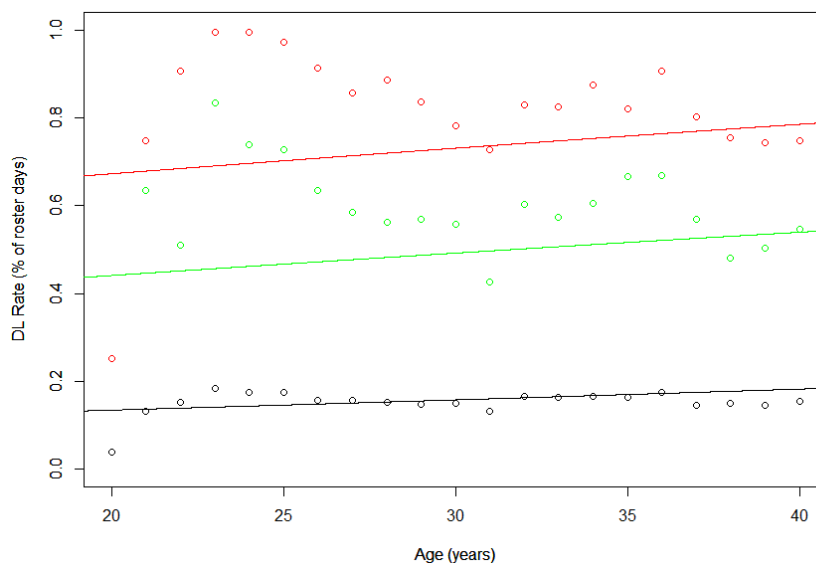|	Estimate	Std. Error	t value	Pr(>|t|)
Intercept	-1.25255	0.01271	-98.51	< 2e-16

Validation of injury model

Figure 1: The relationship between DL rates and player age.

In Figure 1, the lines plot the two-stage model estimates (at the mean, 90th, and 95th 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 more conservative regarding the tails of the injury distribution, but matches the shape of the data quite well overall.

Parameter estimates for the multivariate normal model of player defensive capabilities

Table 8: 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 9: Random effects variance-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

Example output from the multivariate normal model of player defensive capabilities

Table 10: Top 10 defensive players in terms of UZR/150 at each position

Rank	First base	Second base	Third base	Shortstop
1	Mark Kotsay	Chase Utley	Nick Punto	Nick Punto

2	Casey Kotchman	Nick Punto	Evan Longoria	Adam Everett
3	Albert Pujols	Placido Polanco	Adrian Beltre	J.J. Hardy
4	Ike Davis	John McDonald	Craig Counsell	Brendan Ryan
5	Kevin Youkilis	Ben Zobrist	Scott Rolen	Cesar Izturis
6	Daric Barton	Dustin Pedroia	Placido Polanco	Elvis Andrus
7	Chase Utley	Craig Counsell	Adam Everett	John McDonald
8	Andruw Jones	Brandon Phillips	Jack Hannahan	Clint Barmes
9	Mark Teixeira	Mark Ellis	Ryan Zimmerman	Paul Janish
10	Justin Morneau	Juan Uribe	John McDonald	Craig Counsell

Rank	Left field	Center field	Right field
1	Brett Gardner	Franklin Gutierrez	Brett Carroll
2	Tony Campana	Tony Gwynn	Tony Gwynn
3	Nyjer Morgan	Peter Bourjos	Franklin Gutierrez
4	Tony Gwynn	Carlos Gomez	Andres Torres
5	Andres Torres	Brett Carroll	Tony Campana
6	Carlos Gomez	Nyjer Morgan	Nyjer Morgan
7	Scott Cousins	Tony Campana	Carlos Gomez
8	Jacoby Ellsbury	Andres Torres	Ben Zobrist
9	Brett Carroll	Brett Gardner	Ben Revere
10	Reed Johnson	Ben Revere	Scott Cousins

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.