Daniel Cooney- Harvard University

Title: Look Fastball, Adjust Curveball: A Game Theoretic Approach to Batting

Introduction/ Review of Literature

Since Henry Chadwick invented such metrics as batting average and earned run average in the mid-19th century, baseball has been a game in which participants and followers have tried to find objective methods for evaluating the skills of teams and players. Throughout most of the 20th century, teams used simple statistics like home runs and batting average, which appeared to be obvious indicators of a batter's skill, to determine which players would be included on their roster. In recent decades, the field of sabermetrics has emerged, using much more advanced methods of statistical analysis in order to evaluate the skill of major league players. For instance, Michael Lewis notes in his book *Moneyball* that Oakland Athletics' general manager emphasized on-base and slugging percentage in drafting new players (Lewis, 2003). In general, statistics like value over replacement player (VORP) and win shares, which attempt to measure how well individual players contribute to team victories, are gaining momentum in the world of statistical analysis in baseball.

While many of these statistical measures used in baseball were originally designed for the purpose of sports, some measurements of player quality were grounded in previously defined mathematical theories. One example of this is the Markov chain, which originally comes from probability theory, and bears the property that each individual step in a series of events is not affected by the outcome of the previous one. In other words, the Markov property is defined such that these events have no memory. In a Markov chain analysis, there are two components: the state, the number of base runners and outs, represented by a transition matrix T, and the probability that various results occur in a given an at-bat, represented by a vector V. When the transition matrix is raised to large enough of a power, than the entries of TV asymptotically approach a final value that represents the true impact of a player based upon his probability of achieving various results in a given at-bat (Pankin 1987).

However, it is necessary to specifically note that this Markov chain model is based upon the assumption that the result of each at-bat is an independent, identically distributed random variable. It assumes that the result of a batter's given at-bat is not at all effected by the result of the previous at bat, and that it has no bearing on the result of the next at-bat. This is a fair assumption if you want to analyze baseball through Bayesian lens, thinking that, for instance, the number of home runs a player has hit in a season reveals the probability that the player will hit a home run in any given at-bat. However, this is clearly not the case in practice. What is missing from the Markov chain model is the uncertainty the batter faces when preparing to hit against a given pitcher. Batters usually have some guess as to what pitch the pitcher will throw, which will help them to achieve a better result if they guess correctly, and converse, will put them at a disadvantage if they guess wrong. While the result of a given at-bat certainly is affected by the physiological skills of the batter and the pitcher, we must note that both the pitcher and batter can enter the at-bat, and approach any pitch within, with both a memory and a strategy, which means that the batter-pitcher showdown can be modeled using game theory.

The interaction between the pitcher and batter is essentially a repeated game, and can be modeled and analyzed in a manner similar to the way in which the Axelrod tournament modeled and evaluated the iterated prisoners dilemma. Kenneth Kovash and

Steven Levitt analyzed data from Major League Baseball to see whether pitchers used optimal mixed strategies in selecting which pitch to throw. They considered a strategy that follows the minimax principles to be optimal, and they found that the pitchers deviated from this optimal plan in two important ways. The first issue is that pitchers often chose to throw a non-fastball with a much higher probability after having thrown a fastball than after having thrown a similar breaking pitch, which can make their strategy more predictable, and thus less effective, against the batter. The other fault they found was that pitchers through too many fastballs, which again makes the pitcher's strategy much more predictable (Kovash, Levitt 2009). What the paper ultimately shows is that, so far, pitchers do not follow an optimal strategy for selecting pitches to thrown. It does not, however, give a systematic definition of what such a strategy should be, so research is still required to find this strategy.

This project intends to explore different strategies pitchers and hitters in an attempt to see which strategies produce the greatest possible outputs for the players involved. However, because it is much easier to quantify batter's ability for hitting various pitch types than pitcher's ability to throw them based on readily available baseball statistics, the analysis in this project will limit itself to evaluating the success of batters. Additionally, given a limited amount of programming sophistication, this project will only analyze pitching and hitting strategies that do not rely on stochastic properties, so these strategies are unfortunately insufficient in satisfying all of the requirements for a minimax strategy.

Limitations aside, the goal of this project is to test a set of four strategies used by the pitcher and hitter, and to determine whether which of these strategies results in the greatest output for the hitter. To avoid over-complication, the choice of pitches has been limited to fastball, F, and curveball, C.T he four strategies that were considered were as follows:

1. TFT: Analogous to the Tit-for-Tat strategy in the iterated prisoner's dilemma in that a player mimics the opponents last action, the batter will anticipate the pitch that the pitcher threw in the last at-bat.
2. OTFT: Very similar to the TFT strategy, except that the O stands for opposite, meaning that the player, given that the opponent made action "X", will choose action "not X" in the next round. In this case, the batter will anticipate the pitch that the pitcher did not throw in their previous interaction.
3. Best Bet: The best bet strategy is simply that the player will always go with the action that can potentially give them the highest possible output, so a batter will always choose to look for the pitch that will give them the highest output if they, in fact, are thrown that pitch.
4. ATF: this strategy is called adjusting to fifty-fifty, essentially meaning that given the player has a memory of three previous rounds, the player will choose the action that has shown up in the minority throughout those rounds. In this model, the batter would anticipate, for instance, a curveball if two or more of the past three pitches have been fastballs. Also, note that for the first three pitches, for which there is no possible full memory, the player will automatically choose their best pitch to either hit or throw.

These strategies were assumed to be viable options for both the pitcher and hitter. For concreteness, a simulation in many ways similar to the Axelrod tournament was designed in order to determine the best of the given strategies, and this simulation consisted of pitting the nine batters in the Boston Red Sox starting lineup against the five pitchers in the New York Yankees starting rotation. It was hypothesized that one of the two reactionary strategies, TFT or OTFT, would defeat the two strategies, ATF and Best Bet, that do not in any way rely on the strategy of the opponent because a batter using either TFT or OTFT is harder to be identified by a pitcher because the batter's actions are influenced by the pitcher's own strategy. Further it was hypothesized that TFT would outperform OTFT, at least in a world with just the four given strategies, because the two strategies theoretically perform about equally well in a world of TFT, OTFT and ATF, but OTFT would do very poorly against the simple Best Bet strategy because it deliberately avoids the last pitch thrown, which will always also be the next pitch thrown for a Best Bet Pitcher.

In addition to the four strategies mentioned above, a batter strategy that incorporated both an interpretation of the batter's strengths and the pitcher's weaknesses were also explored, albeit less rigorously. The strategy was named FBP, which stands for fifty-fifty batter pitcher, which means that for any given at-bat, the batter essentially flips a fair coin, so with probability ½ to anticipate the throwing of the pitcher's best pitch to throw and with probability ½ to anticipate seeing the batter's own worst pitch to hit. Checking with batters chosen as representative as fastball and curveball hitters, the FBP strategy can be compared as a batting strategy in the same manner as in the primary experiment. It was hypothesized that FBP would be outperformed by the best performing strategy from the prior experiment because although it takes into account properties of both the pitcher and the batter, it has no reactionary element to it that can be used to interpret revealed information about the pitcher to be applied in later at-bats.

However, it is important in general to consider such strategies that involve both information about the pitcher's and batter's qualities because they certainly can easily be applied to baseball situations. A pitcher whose best pitch is a fastball, when encountering a batter whose best pitch to hit is also a fastball, may very well be best served to throw a curveball to avoid the batter receiving his optimal output, even though the pitcher would receive his best output with the fastball if the batter ultimately anticipated a different pitch. Even more important, it is necessary to note that considering such a strategy has a greater bearing in general to game theory than just its application to pitch selection in baseball. In any situation, when a decision maker must choose between focusing on his own strengths and his opponents weakness must consider the tradeoff that the batter faces when he employs a strategy like FBP. So ultimately, the relevance of this study of game theory in baseball does not solely impact the strategy of players, coaches and baseball officials, but also has implications in general for the study of repeated games.

Methods

A repeated game theory tournament, somewhat similar to Axelrod's tournament, was designed. The general structure was that the eighteen batters on the 2011 Boston Red Sox and 2011 New York Yankees opening day lineups were assigned outputs for the four possible situations: being thrown a fastball when expecting a fastball, being thrown a curveball when expecting a curveball, and being throw one of the two pitches when expecting the other of the two pitches. These outputs can be expressed in a typical game theory payoff matrix, in which J_K is the output generated when pitch J is thrown and the batter anticipated pitch K, which is shown below. For the outputs, an arbitrary assumption was made that if a batter was thrown the pitch, A, of the two for which he normally receives a higher output, but he was not expecting it, then his output A_B would be equal to ½ A_A. Then, if this batter were thrown the pitch for which he were a weaker hitter, B, he would then have output B_A=(1/3) B_B. Although the assumption is arbitrary that a batter gets one-half of his ideal output for hitting his stronger pitch when surprised by the pitch, while only one-third when it is the batter's weaker pitch, it makes intuitive sense that a batter should be able to reflexively hit his better pitch than his weaker pitch when surprise, if he is in general better at hitting the given stronger pitch. So below, there is also a payoff matrix included that illustrates this rule that has been defined for the relation between output for pitches expected against those faced when the other pitch was anticipated.

General Batter	Fastball Thrown	Curveball Thrown
Fastball Anticipated	F_F	C_F
Curveball Anticipated	F_C	C_C

Batter who is more adept with fastballs	Fastball Thrown	Curveball Thrown
Fastball Anticipated	F_F	(1/3) C_C
Curveball Anticipated	(1/2) F_F	C_C

Batter who is more adept with curveballs	Fastball Thrown	Curveball Thrown
Fastball Anticipated	F_F	(1/2) C_C
Curveball Anticipated	(1/3) F_F	C_C

These outputs were derived from the closest readily available statistic that bore relevance to a batter's ability to hit different types of pitches. It was assumed that a batter's batting average against pitcher's classified as fly ball pitchers corresponds to the batter's ability to hit against fastballs, because fastballs fall less in their path to the plate than curveballs, so the batter is more likely to impart backspin onto the ball and to hit it in the air, so a fastball pitcher would tend to also be a fly ball pitcher. We then made the similar assumption that a batter's ability to hit curveballs was correlated to his batting average against pitchers classified as groundball pitchers, as curveballs tend to fall significantly in their path to the batter, so the batter is more likely to impart topspin on the ball, so curveball pitchers are likely to be groundball pitchers. Because the 2011 rosters were being examined, the 2010 players batting averages against ground and fly ball pitchers were collected from Baseball-Reference.com, except in the case of players injured in 2010, for whom data was taken from 2009.

Having collected the batting average for each of the Red Sox and Yankees, the average was then taken for both fly balls and ground balls. The average batting average for fly balls was defined to have output of 10, so the average player on the two teams would have output of 10 when hitting a successfully anticipated fastball. This system of measuring output was arbitrarily scaled such that for every 10 points in batting average a player differed from the two team mean, his output for the given pitch differed by 1 from the home point of 10. Below is a chart of all of the bating averages of the Red Sox and Yankees, as well as their corresponding outputs (batting averages from baseball-refernece.com).

Name	Fly Ball Batting Avg.	Ground Ball Batting Avg.	Output for Fastball	Output for Curveball
Ellisbury	0.279	0.282	10.67777778	10.97777778
Pedroia	0.276	0.229	10.37777778	5.677777778
Crawford	0.258	0.323	8.577777778	15.07777778
Youkilis	0.288	0.317	11.57777778	14.47777778
Gonzalez	0.335	0.293	16.27777778	12.07777778
Ortiz	0.257	0.284	8.477777778	11.17777778
Cameron	0.225	0.3	5.277777778	12.77777778
Saltalamacchia	0.315	0.206	14.27777778	3.377777778
Scutaro	0.276	0.256	10.37777778	8.377777778
Martin	0.214	0.278	4.177777778	10.57777778
Teixiera	0.263	0.279	9.077777778	10.67777778
Cano	0.324	0.304	15.17777778	13.17777778
Jeter	0.339	0.214	16.67777778	4.177777778
Rodriguez	0.271	0.258	9.877777778	8.577777778
Gardner	0.257	0.254	8.477777778	8.177777778
Jones	0.236	0.211	6.377777778	3.877777778
Swisher	0.267	0.305	9.477777778	13.27777778
Posada	0.22	0.284	4.777777778	11.17777778
Average	0.272	0.270	10	9.87

Now having the outputs of the hitters, it was needed to find some method for determining pitcher's tendencies for throwing fastballs and curveballs. Essentially, the metric used was the ratio of fly balls allowed by a pitcher compared to the total umber of gro8ndballs and fly balls allowed by the pitcher. These statistics were taken from the 2010 numbers of were the members of both the Yankees and Red Sox starting rotations according to

Baseball-Reference. Pitchers whose ratio was greater than 0.5 were deemed fly ball, and thus fastball, pitchers, while lose with fewer than 0.5 were considered curveball pitchers. Below is a chart detailing the number of fly balls and groundballs allowed, the calculated ratio and the resulting classification of the pitcher (number of fly balls and groundballs courtesy of bsesball-reference.com).

Pitcher	Fly Balls Allowed	Ground Balls Allowed	Ratio: # Fly Balls/(# Groundballs + # Fly Balls).	Classification: Curveball or Fastball
Lester	142	287	0.331002331	Curveball
Lackey	226	310	0.421641791	Curveball
Buchholz	154	259	0.372881356	Curveball
Beckett	136	185	0.423676012	Curveball
Matsuzaka	190	149	0.560471976	Fastball
Sabathia	232	346	0.401384083	Curveball
Burnett	209	256	0.449462366	Curveball
Hughes	98	79	0.553672316	Fastball
Colon	85	92	0.480225989	Curveball
Garcia	193	213	0.475369458	Curveball
Average	166.5	217.6	0.45	Curveball

With both outputs for the batters and classification of the pitchers completed, it was possible to begin simulating a series of potential seasons between the Red Sox and Yankees. Because of computational limitations, our simulation consisted of the batters for the Red Sox playing against the pitchers for the Yankees. In simulating the season, it was assumed that each batter faced each pitcher a random number of times throughout the season. Because the Yankees and Red Sox compete eighteen times in a season in six series of three, it is reasonably safe to bound the number of possible at-bats between 5 and 25. The number of at-bats between a given batter and pitcher was chosen by picking a random integer within this range.

After assigning each matchup between batter and pitcher a number of at-bats, a strategy was selected from TFT, OTFT, Best Bet and ATF, and a strategy was similarly assigned to the batter. Each of the four strategies for the batter were tested against each of the four pitcher's strategies, for a total testing of sixteen different strategy mixes for each batter-pitcher showdown. It should be noted that for all of the strategies except ATF, the batter and pitcher as a rule choose their best pitch for the first at-bat, while ATF involved the player using his best pitch for the first three at-bats. For the simulation, each at-bat was assumed to last for a single choice of pitch, and the result of each at-bat was the output chosen appropriately from the payoff matrix. For both the whole Red Sox lineup and each individual player, the total output was divided by the number of at-bats to determine the batter's output per at-bat for a given strategy, which is a measurement of the strategy's effectiveness.

The strategy using both the batter's strengths and pitcher's weaknesses was also tested, but less extensively. Taking Adrian Gonzalez and Kevin Youkilis to serve as representative fastball and curveball pitchers, respectively, the FBP strategy was tested against the strategy deemed most effective from the above simulation. The rest of the simulation was conducted in a manner identical to the above experiment.

Results

The season was simulated for the 16 possible matchups of pitcher and batter strategy. Both the total output of the batter and the output per single at-bat were calculated. The simulation consisted of 661 at-bats, distributed randomly between the different matchups between batter and pitcher. The strategy matchup with the highest overall output, and thus also output per at-bat, for the entire team of Red Sox was the batter employing the Best Bet strategy against pitcher employing TFT, which results from the pitcher immediately switching to imitate the batter, who does not change his strategy, allowing the batter to see the pitch he expected from at least the second at-bat until the end of the interaction between the given pitcher and batter.

A table of results for each of the strategies is shown below. The best overall batter strategy was TFT, which had total output of 22,393.59 over the 661 at-bats, for an average payoff per at-bat of 8.47. In descending order, the other three strategies were OTFT, ATF and Best Bet, which had respective outputs per at-bat of 7.73, 7.51, and 7.03. This confirmed both of our hypotheses that TFT and OTFT would outperform Best Bet and ATF and that of the first two, TFT would outperform OTFT. This second, and less obvious result, can be attributed to the fact that against TFT, OTFT and ATF, the two strategies each received various outputs per at-bat in the range of approximately 7 to 8. However, against a pitcher using the Best Bet strategy, the batters using TFT received an output per at-bat of 10.17, while the batters using OTFT only received an output per at-bat of 5.68. Another detail emerging from the results was that the ATF strategy outperformed the Best Bet strategy. While no hypothesis was offered on distinguishing between the two strategies either way, this result is certainly consistent with the idea that a batter using a more varied strategy prevents the pitcher from figuring out the strategy and exploiting the strategies weaknesses.

Batter Strategy	Pitcher Strategy	Output	Output per At-Bat	Average Output per At-Bat
TFT	TFT	5697.762	8.619172909	
TFT	OTFT	4807.92	7.273082627	
TFT	ATF	5165.79	7.814443148	
TFT	Best Bet	6722.12	10.16874952	**8.468862051**
OTFT	TFT	4937.22	7.468678557	
OTFT	OTFT	5061.3	7.656378038	
OTFT	STF	4829.56	7.305818094	
OTFT	Best Bet	3756.57	5.682674421	7.028387278
ATF	TFT	5238.26	7.924070658	
ATF	OTFT	4493.01	6.796708966	
ATF	ATF	4896.52	7.407110464	
ATF	Best Bet	5235.626	7.920086129	7.511994054
Best Bet	TFT	8242.84	*12.46918759*	
Best Bet	OTFT	2030.44	3.071506574	
Best Bet	ATF	5061.46	7.656620075	
Best Bet	Best Bet	5107.44	7.72617538	5110.545

In addition to the results of the entire group, the best strategy for each of the individual hitters was calculated in the same manner, using total output and output per at-bat. Out of the (Red Sox batters, 6 achieved the best output using TFT, while 2 found OTFT best and 1 was most successful with the Best Bet strategy. The team consisted of 4 fastball hitters and 5 curveball hitters, of whom 3 out of 4 and 3 out of 5 found TFT to be their best strategy. This implies that the distribution of best strategies was not heavily dependent on the batter's relative strength between the two possible pitches. This most likely suggests that the greatest contributing factor, at least over the short simulation of one season, was the randomized number of at-bats that each batter had against each pitcher, which may have skewed the way in which batter pitch preference affects the effectiveness of a batter's strategy.

Player	Average Output per At-Bat	Best Strategy	Output per At-Bat for Best Strategy
Ellisbury	8.59	TFT	9.29
Pedroia	5.44	OTFT	5.78
Ortiz	8	TFT	8.42
Youkilis	10	TFT	11.35
Gonzalez	10	TFT	10.57
Crawford	9	TFT	11.49
Cameron	8.52	OTFT	9.22
Saltalamacchia	5.1	Best Bet	5.75
Scutaro	6	TFT	6.71

After determining that TFT was the best of the four strategies originally tested, a new simulation was run in which representative fastball and curveball hitters Adrian Gonzalez and Kevin Youkilis, respectively, were once again run through simulated seasons, but in this case, the potential strategies were drawn from just TFT and the new FBP. Overall, TFT outperformed FBP with an average output per at-bat of 10.12 for TFT, as compared to the average output per at-bat of 9.44 for FBP. In addition, it should be noted that TFT outperformed FBP in both cases in which the pitcher followed TFT and when the pitcher followed FBP: TFT batters achieved 9.52 to FBP's 9.39 when competing against TFT pitcher, and the TFT batters achieve 10.74 to FBP's 9.48. This means that, in the situation that has been posed, TFT is a dominant strategy over FBP.

Batter	Pitcher	At-Bats	FBP vs TFT	TFT vs FBP	FBP vs FBP	TFT vs TFT
Youkilis	Sabathia	5	54.57	40.97	37.15	72.5
Youkilis	Burnett	22	174.06	231.28	225.2	319
Youkilis	Hughes	21	203.08	180.82	216.28	111.54
Youkilis	Colon	11	98.46	123.59	119.21	159.5
Youkilis	Garcia	16	153.44	126.86	122.89	232
Gonzalez	Sabathia	12	133.83	129.68	125.53	72.8
Gonzalez	Colon	8	61.88	72.98	77.07	97.1
Gonzalez	Hughes	7	72.38	72.89	64.93	113.96
Gonzalez	Burnett	13	145.14	117.43	105.36	80.94
Gonzalez	Garcia	5	43.64	48.69	36.37	32.41
Total	Total	120	1140.48	1145.19	1129.99	1291.75
Average per at bat			9.47910156	9.518248733	9.391913906	10.73638244

Discussion

Although three hypotheses were tested and confirmed as a result of this experimental model, it is important to note that these successful test do not necessarily correlate to immediately correlate to successful implementation into actual gameplay. This discussion will describe potential weaknesses and sources of error in the model and simulation, but more important, it will focus on ways in which research on this game theoretic model of batting can be enhanced and generalized so that, sometime in the future, this research could potentially lead to changes in coach and player strategy.

The simulations that were completed have several potential sources of error, which have arisen from the choice to use the particular data set of the 2011 Yankees and Red Sox. Two of the potential sources of error actually have the exact opposite effects and implications. Drawing from just two teams in one year of baseball history, there are only eighteen data points for hitters and ten for starting pitchers, which may mean that the sample is either not large enough or nor sufficiently varied to give a true insight into the impact of batters' relative abilities for hitting fastballs and curveballs have on the effectiveness of various batting strategies. On the other hand, while the average batting averages for groundballs and fly balls, as well as the average ratio of fly balls to total groundballs and fly balls, were based of the entire set of Yankees and Red Sox hitters and pitchers, while the experiment was only run on Red Sox hitters and Yankees pitchers. This was caused by the fact that it had been originally intended for both the offense and defense of each team to be tested, but that, after comparing the Red Sox offense to he Yankees pitching, it appeared that completing the reciprocal simulations by the methods used was computationally and intellectually inefficient. However, the simulations that were completed had been based off of levels of output indexed to average statistics for both the Yankees and Red Sox, which may mean that some of the results that may potentially have implications about particular "fastball hitters" or "curveball hitters" on the Red Sox that do not hold true if our models were simply drawn from the Red Sox sample. However, it is likely that these effects contributed little to conceptual error, because it is the relative value of the batter's ability to hit the two pitches that appears to matter more than the batter's absolute ability, and a player' comparative ability to hit fastballs and curveballs is much less affected by changing the index or scaling. Now that the actual potential errors in the running of the simulation have been addressed, it is important to evaluate the model in the context of its ability to accurately relate to an actual baseball game.

One issue with the model at hand is that it treats at-bats as the smallest unit on which the batter and pitcher can implement their strategy, while much of the actual game theory and decision-making in baseball is made throughout the at-bat, observing how one's opponent acts on the first pitch of an at-bat has an effect on the decisions the players will make for subsequent pitches, while our model completely ignores that. Kovash and Levitt made a similar omission of theory within the at-bat, focusing only on the choice of and result from the terminal pitch of an at bat. This assumption has drawn criticism from writer Phil Birnbaum, who argues that focusing on the terminal pitch of an at-bat in some ways trivializes the value of throwing fastballs (Birnbaum 2009). This seems like a logical argument: pitchers are typically more accurate with their fastball, so they are more likely to throw a fastball early in the count with the hope of getting a strike on the batter. On the other hand, breaking pitches like curveballs are often thrown later in the count, and thus have a higher probability of being a terminal pitch because of either the occurrence of a walk or strikeout, or additionally the batter's increasing likelihood to take less of a risk with his swing in order guarantee making contact with the pitch. While this assumption that pitchers and batters choose to throw and look for just one pitch in an at-bat, much of the significant game theory, as well as the practicability of comparing the mathematical exercise with the actual result in baseball, becomes lost. In order to truly model baseball through game theory, one cannot just pare the game down to results on terminal pitches, but must rather model the entire interaction between pitcher and hitter that occurs throughout the at-bat.

Another issue with simply accepting a baseball game as simply one that fits neatly into the construct of a repeated game is that it's results are much more complex that those of the prisoner's dilemma or a game of chicken. While, for instance, in a game of chicken an output of 10 could easily be defined as the glory received from driving past the swerving opponent and -20 the physical and financial harm imparted by engaging in a car crash, it is much harder to explain what our defined 10 points of payoff means in the context of baseball. The actual results of a batted ball are to complicated to model using a simple 2 x 2 matrix used in a typical game theory problem, so it may be the Markov chain model that has appeared inapplicable to real life due to its lack of game theory that ultimately serves to help the oversimplified game theoretic model achieve a better grounding in reality. If someone were to complete a Markovian analysis to the value of various batting results given the pitch type, and in some way were able to figure out or account for the intended pitch type, then, for instance, one would be able to know that value and probability of hitting a double when a given batter anticipates a fastball against a given pitcher and is then thrown a fastball. Knowing the probabilities and values of all the given states of anticipating and seeing pitches would allow a researcher to test various strategies using mathematical simulations, which would carry out both the execution of the strategy and then also a simulation of the baseball result given the strategies employed by the pitcher and the hitter. The results of such simulations could potentially demonstrate which strategy is necessary to achieve the maximum chance of obtaining the maximum possible output, satisfying Kovash and Levitt's minimax conditions.

Another potential negligent assumption made in the creation of this model is that the discount factor, w, from round to round is simply 1, meaning that each at-bat has equal potential for impact with respect to the total payout. While w is generally either treated as 1 because the payoff of the last round is constructed to have no impact on future payoff, or is defined to be less than 1 because it is constructed to have greater benefit, for instance, to choose to defect and hurt someone early on than it is to defect and hurt the person later on. In baseball, however, it is possible that there exists w that varies dependent upon the payoffs of immediately prior rounds. For instance, if a batter were to strike out in a given at-bat, he may be more likely to be conservative in his next at-bat in order to guarantee making contact, which would decrease his maximum payout for hitting, and, as a result, w would be decreased for that next at-bat. At the same time, a player doing poorly in a given at-bat may inspire the batter or give the batter some insight into the pitcher that would allow him to be more likely to do well in the next at-bat, which corresponds to an increase in w. It is possible, and could be the subject of future research, that w for baseball follows a biased random walk. That is to say that if the batter, for example, succeeds in a given at-bat, perhaps anticipating the correct pitch, he has some probability that his w will moderately increase, a smaller probability that w will moderately decrease, and even smaller probabilities that w will increase or decrease dramatically. Such a random walk may be partially responsible for the existence of slumps and hitting streaks, and the study of whether this weighting factor w follows and random walk would provide significant insight into whether such hot and cold streaks are cause by physical or psychological aspects of the batter's approach to the game.

Conclusion

The goal of this project was to see if repeated game theory could be used to evaluate the effectiveness of various strategies for a pitcher throwing a fastball or curveball, and for the batter anticipating to hit one of the two pitches. A game theoretic model was developed that was supposed to convert statistics indicating batter's abilities to hit fastballs and curveballs into a payoff matrix, which was then used to evaluate which of four defined pitching and hitting strategies, TFT, OTFT, ATF, and Best Bet, worked best when a hitter is pitted against pitchers using these strategies. The simulations showed that TFT was the most successful strategy of the four, largely due to the fact that it is a responsive strategy in that it allows the batter to respond to changes in the pitcher's actions. However, because a batter in baseball is typically aware of both his own strengths and weaknesses, but also those of the pitcher, it would appear important to have a strategy that considers both the batter's strengths and the pitcher's weaknesses. Although it was computationally too tedious to consider a strategy like this in detail without reasonably sophisticated computer programming, one such strategy, FBP, was run against the best of the strategies that considered only the batter's strengths, TFT. TFT proved to be a dominant strategy over FBP, but that may have been caused largely by the fact that FBP is in itself a rather simple strategy that does not actually respond to the opponent, so the batter could very well employ the strategy without noticing necessary information about the pitcher. FBP still did relatively well with respect to TFT, which means that it is still important to have a strategy that acknowledges both the pitcher and hitter's skills. Thus it appears clear than an optimal strategy for anticipating pitches as a batter would be some form of combination of TFT and FBP's stronger qualities, a strategy that both responds to the pitcher's actions but also takes into account the batter's strengths and weaknesses. It is probably impossible to have a strategy that can always take care of both concerns in a given at-bat, which would mean that there would need to be some sort of probabilistic element at the at-bat level to achieve a beneficial combination of TFT and FBP, which would satisfy one of Kovash and Levitt's previously unsatisfied conditions for a minimax strategy that achieves optimum payoff for the hitter over time. It is apparent that one can also use other sabermetric methods described in the introduction, such as win probability contribution or Markov chains, in order to enhance the value of this game theoretic model to its actual application in baseball, which clearly is one of the ultimate goals of future research.

Literature Cited

Birnbaum, Phil. "A Game Theory Study on Pitch Selection." *Sabermetric Research.* 27 Sept. 2009. Web. 12 May 2011. <http://sabermetricresearch.blogspot.com/2009/09/game-theory-study-on-pitch-selection.html>.

Kovash, Kenneth and Levitt, Steven D., Professionals Do Not Play Minimax: Evidence from Major League Baseball and the National Football League (September 2009). NBER Working Paper Series, Vol. w15347, pp. -, 2009. Available at SSRN: http://ssrn.com/abstract=1472285

Lewis, Michael. *Moneyball: the Art of Winning an Unfair Game.* New York: W.W. Norton, 2003. Print.

Pankin, Mark. "Markov Baseball Models Theory." *The Great American Baseball Stat Book* (1987): 520-24. *Mark Pankin Home Page.* Web. 12 May 2011. <http://www.pankin.com/markov/theory.htm>.