The Theory of Theft:*

An Inspection Game Model of the Stolen Base Play in Baseball

by Theodore L. Turocy

Department of Economics

Texas A&M University

College Station TX 77843

Phone: 979.862.8082

Fax: 979.847.8757

Email: turocy@econmail.tamu.edu

December 17, 2004

Abstract

This paper applies the theory of equilibrium in mixed strategies in an inspection game model to describe the strategic interaction in the stolen base play in baseball. A parsimonious simultaneous-move game model offers predictions about how the observable conduct of the teams on offense and defense responds as the characteristics of the players involved change. The theory organizes observations from play-by-play data from Major League Baseball, where highly-motivated, experienced professionals interact in an environment where private information is not significant.

JEL Classification Numbers: C72, C73

Keywords: mixed strategy, baseball

1 Introduction

The suggestion of using formal techniques to study questions of strategy in sport dates back in print at least as far as Mottley [4]. This suggestion was fulfilled, in the case of the game of baseball, in part by the work of Lindsey [3] and Bellman [1], who used ideas from basic decision theory and dynamic programming, respectively, to investigate some questions of optimal strategy.

*. This document has been written using the GNU $\TeX_{\text{MACS}}$ text editor (see www.texmacs.org).

Some more recent papers have departed from this normative approach towards a descriptive view. In particular, the zero-sum nature of many sporting contests makes it likely that active randomization is part of optimal strategy in situations where opponents make moves that are essentially simultaneous. Walker and Wooders [7] compared predictions from the theory of two-player simultaneous-move zero-sum games to observations of service behavior in championship tennis matches. In that setting, the authors cannot reject the hypothesis that servers win the same number of points when serving to the opponent's forehand versus their backhand, as predicted by the theory. Additionally, they find that professionals seem to do better at choosing their behavior in a serially uncorrelated fashion, compared to behavior reported in laboratory games. This was good news for the theory, as it has generally not done well in describing behavior of subjects in comparable laboratory games, and can be interpreted as indicating that motivated experts may in fact conform more closely to the predictions of theory.

Chiappori et al [2], like Walker and Wooders, take mixed strategies as a starting point in investigating behavior in penalty kicks in professional soccer. Unlike the repeated interaction of tennis, a given striker and goaltender will face each other in a penalty kick situation at most a few times in a season. Therefore, the authors focus on hypotheses about aggregate behavior which are robust to the introduction of heterogeneity across players. Mixed-strategy equilibrium theory predicts some empirical regularities that make sense to the seasoned soccer observer: for example, that right-footed strikers should kick towards their right more often than goalies will dive in that direction. The authors show evidence that the theory of mixed-strategy equilibrium again organizes the data well.

This paper turns this theory to the analysis of the stolen base play in baseball. In parallel with the terminology used in baseball, the stolen base play is conceptualized as a simple inspection game. In the game, the defense chooses between "inspecting," which gives the defense a positive probability to cause the stolen base play to fail, and not inspecting, in which case the stolen base play, if attempted, is successful for the offense. Heterogeneity in skill across players maps in the model to different levels of effectiveness of the inspection "technology." It is shown that linear relationship exists in equilibrium between the probability the stolen base play is attempted and the probability the play is successful.

The observed dataset for the stolen base play shares characteristics with that of the penalty kick in soccer. While situations in the play of a game of baseball where a stolen base can be attempted occur relatively frequently, the same individual players do not participate against each other in these "stage games" repeatedly. This contrasts with the repeated interaction of tennis players within a match which Walker and Wooders were able to exploit. However, individual baseball players participate in these stage games much more often than soccer players participate in penalty kicks over the course of a season. This makes it feasible to investigate predictions of the theory relating to heterogeneity, which was not possible for Chiappori et al.

The paper is organized as follows. Section 2 outlines the stolen base play in baseball and motivates the design of the model. A parsimonious idealization of the stolen base game as an inspection game is developed in Section 3. In the model, the heterogeneous abilities of players play the role of varying the efficacy of an inspection procedure. While these abilities are not observable, the model does have comparative statics predictions which can be expressed in terms of observable quantities. These predictions are taken to detailed data from two decades of Major League Baseball games in Sections 4 and 5. Section 6 concludes.

2 A primer on the stolen base play

In baseball, two competing teams take turns on offense and defense. While on offense, each team attempts to score runs by advancing team members around a sequence of four bases. The team's turn on offense is terminated when three of its members have been "put out," which can occur by various means; therefore, outs are a scarce resource for a team. The game is won by the team scoring the most runs after nine innings (i.e., nine turns on offense for both teams); ties are broken by playing successive extra innings.

The players on a baseball team bat in a strict rotation (the "batting order"). As in many bat-and-ball games (such as cricket or rounders), most advancement occurs on batted balls, that is, when the current batter successfully strikes a pitched ball. However, attempts to advance in baseball are permitted at any time. An attempt to advance a base without the benefit of the ball being batted is called an attempt to "steal" a base. A player successful in stealing a base has advanced one base closer to his ultimate goal of scoring a run; a player who is unsuccessful is put out, costing the team one of its scarce outs.

A stolen base attempt can be thought of as a race. In this race, the offensive player (the runner) runs a distance of about 90 feet. His opponents are members of the defense, the pitcher and the catcher. The pitcher pitches the ball to the catcher, a distance of about 60 feet, and then the catcher relays the ball to the base, a throw of an additional 120 feet. If the runner reaches the base prior to the relay throw from the catcher, the attempt is succesful; if not, the attempt fails. While the runner need not wait until the pitcher starts to pitch the ball to begin running, doing so would result in almost certain failure, as the pitcher is permitted instead to throw directly to the base to which the runner is advancing. Therefore, the runner wants to time his departure to closely match the start of the pitcher's throw to the plate. The choices of the offense and defense are made essentially in ignorance of each other, making a simultaneous-move model the natural choice for describing the interaction.

The interaction between a pitcher and a player known for stealing bases is often described in language suggestive of mixed-strategy equilibrium. Pitchers are encouraged to make the runner's ability to time his departure more difficult by varying the type, timing, and style of their delivery of pitches; runners, for their part, try to avoid patterns in their behavior, such as always attempting to steal at the first opportunity. Professionals thus perceive some advantage to unpredictability in this setting, which suggests that an equilibrium in mixed strategies should be a feature of an appropriate organizing theory. Finally, the wide availability of data in modern professional sport ensures that common knowledge of the abilities of the relevant players obtains; therefore, it is reasonable to model this interaction as one of perfect information regarding the revelant parameters.

3 The stolen base play as an inspection game

A particular state of a baseball game can be described by a state vector containing, for example, the inning, score, number of outs, and other relevant factors. Suppose, as in the previous section, that the offense has a runner on first base. At such a point in the progress of the game, the future continuations of the game can be summarized by a vector v, which expresses the probability the team currently on offense will eventually win the game, conditional on the outcome of the interaction to be described next.

To focus on the stolen base play, the interaction is modeled as a simultaneous-move game. In this game, the offense chooses whether to attempt the stolen base play (strategy S) or not (strategy N). Meanwhile, the defense chooses whether to focus their efforts on trying to put out the batter currently at the plate (strategy B) or trying to interdict a possible stolen base attempt by focusing on the runner (strategy R). If the stolen base play is attempted, there are two possible outcomes: success, resulting in the runner reaching the next base safely, and failure, resulting in the runner being put out.

	batter (B)	runner (R)
attempt (S)	v_S	$\rho v_S + (1-\rho)v_F$
no attempt (N)	v_B	v_R

Table 1. A model of the stolen base play as a zero-sum game between the offense and the defense. In the table, the offense is the row chooser, and the defense the column chooser. The cell entries are the payoffs to the offense, measured as the probability the game will eventually be won by the team on offense.

The structure of this game is presented in Table 1. It is assumed that both teams seek to maximize the probability of eventually winning the game. Each entry in the table is the probability the team on offense will eventually win the game, conditional on the corresponding strategy profile being chosen. The payoff to the defense is one minus that of the offense.

The vector v of continuation values has four components. The continuation value after a successful attempt is v_S, and v_F is the continuation value after a failed attempt. The continuation values v_B and v_R describe the continuations where the play is not attempted.

The defense's strategy R is their "inspection" strategy. When playing this strategy, only a fraction ρ of stolen base attempts are successful. An attempt is always successful when the defense does not inspect and plays strategy B.

Four inequalities are assumed to hold, which jointly ensure the equilibrium is unique and involves active randomization by both sides.

- $v_S > \rho v_S + (1-\rho)v_F$. This will hold if $\rho < 1$ and $v_S > v_F$. This states that a team on offense prefers success to failure, which must be true since no team has a strategic incentive to have the runner deliberately put out in this setting.

- $v_S > v_B$. This says that a succesful attempt is better than just letting the batter hit the ball. This will hold since after a successful attempt, the batter's turn continues, while the runner has advanced towards scoring a run.

- $v_R > v_B$. This encodes an assumption that inspection is costly in the event that an attempt does not occur. This cost derives from the idea that the defense must modify their pitching approach to the batter, and is motivated by the following logic. If there were no runner on base, the pitcher and batter can be thought of as engaging in their own zero-sum game (not explicitly modeled here), with the pitcher choosing the type and location of pitch to throw, and the batter forming expectations about the pitch. Now, with the baserunner on first base, let the strategy B in the game corresponds to following the optimal strategy against the batter as if there were no runner on, and R corresponds to following some modified pitching strategy to defend against the runner. In the latter case, it must be that the defense is no longer using their minimax strategy against the batter, and so the batter will perform better. That this effect is important in analyzing this interaction is pointed out by SMITH [5]:

> [A]ny consideration of base running must include indirect and often subtle effects.... The greatest of these indirect effects is of course the intangible of upsetting the pitcher by diverting his attention from the batter.

- $\rho v_S + (1-\rho) v_F < v_R$. This can be rearranged to read

$$\rho < \frac{v_R - v_F}{v_S - v_F}. \tag{1}$$

  This condition says that the runner is not so skilled (or the defensive players so unskilled) that it is always a best reply to attempt.

The ratio appearing in equation (1) is related to the "breakeven" percentages which appear in the analyses LINDSEY [3] and BELLMAN [1]. Those papers view the stolen base as a decision problem, with only the offense making a choice (whether to attempt or not to attempt). Those models only predict that the frequency with which the stolen base play is successful will exceed the critical breakeven percentage, but are silent on the question of optimal attempt frequencies; the mixed strategy equilibrium of this game model will provide sharper predictions.

When the equilibrium is in mixed strategies, the equations for the equilibrium probabilities of attempting the stolen base play, $p^\star_S$, and for focusing attention on the batter, $p^\star_B$, are

$$p^\star_S(\rho) = \frac{v_R - v_B}{(1-\rho)(v_S - v_F) + (v_R - v_B)} \tag{2}$$

$$p^\star_B(\rho) = \frac{v_R - [\rho v_S + (1-\rho) v_F]}{(1-\rho)(v_S - v_F) + (v_R - v_B)} \tag{3}$$

Note that $p_S^\star(\rho)$ is increasing in ρ and that $p_B^\star(\rho)$ is decreasing in ρ: runners who are better at the stolen base play attempt it more frequently, and the defense "inspects" more often as the ability of the runner increases.

In the model, it is assumed that an attempt is always successful when the defense plays strategy B, that is, when the defense does not "inspect." While this is a simplification, it is a plausible approximation. In order to play reach the highest level of the sport, a player cannot be too slow afoot. Furthermore, Section 4 takes this model to a dataset consisting of players for whom stolen base attempts are a salient activity, and among this subset of players it is certain that most, if not all of them, would win the relay race described in Section 2 with a very high frequency if they did not have to worry about the possibility of the R strategy being played.

Anecdotally, what separates the runners successful at stealing bases from those who are not known for stealing bases is their performance when being watched closely by the defense. For example, Oakland Athletics coach Ron Washington has been quoted as saying:

> "A base stealer is a guy who when everyone in the ... yard know he gonna get the bag, he gets the bag."[1]

Interpreted probabilistically, this is a feature of this model. Suppose the defense is expecting the stolen base play to be attempted and therefore plays the inspection strategy R. A talented base stealer will be successful with relatively high probability (i.e., ρ is large when compared to other players).

4 An observable relationship

Since the parameter ρ is not observable, it is desireable to seek relationships between observable quantities. Let $\pi^\star(\rho) \equiv p_B^\star(\rho) + \rho[1 - p_B^\star(\rho)]$ denote the percentage of stolen base attempts which are successful in equilibrium.

Proposition 1. *There is an affine relationship between the frequency of attempts $p_S^\star(\rho)$ and the frequency of success $\pi^\star(\rho)$ in equilibrium as a function of ρ.*

1. Quoted in *Moneyball: The Art of Winning an Unfair Game*, by Michael Lewis, W. W. Norton & Company, 2003, page 265.

Proof. It is asserted that there are constants A and B satisfying

$$\pi^\star(\rho) = Ap_S^\star(\rho) + B.$$

In view of (3), this can be written

$$(v_R - v_F) - \rho(v_B - v_F) = A(v_R - v_B) + B[(1-\rho)(v_S - v_F) + (v_R - v_B)].$$

Collecting coefficients of ρ:

$$\begin{aligned} -(v_B - v_F) &= -B(v_S - v_F) \\ B &= \frac{v_B - v_F}{v_S - v_F}. \end{aligned}$$

Collecting constants:

$$\begin{aligned} v_R - v_F &= A(v_R - v_B) + B(v_S - v_F) + B(v_R - v_B) \\ &= A(v_R - v_B) + (v_B - v_F) + B(v_R - v_B) \\ v_R - v_B &= A(v_R - v_B) + B(v_R - v_B) \\ 1 - B &= A. \end{aligned}$$

Therefore, $p_S^\star(\rho)$ and $\pi^\star(\rho)$ are related according to the equation

$$\pi^\star(\rho) = \left[1 - \frac{v_B - v_F}{v_S - v_F}\right] p_S^\star(\rho) + \frac{v_B - v_F}{v_S - v_F}. \tag{4}$$

□

Since the vector v of continuation values depends on the state of the game, the value of the ratio $(v_B - v_F)/(v_S - v_F)$ also depends on the state of the game. Therefore, the analysis specializes to the case where the stolen base game arises at the beginning of the baseball game. In particular, attention is restricted to situations in which a team's first hitter reaches first base in his team's first turn at bat of the game with a tie score. It is commonplace for Major League Baseball teams to place a player who is considered a thread to attempt the stolen base play in this first place in their batting order; therefore, the stolen base play is a salient part of game at these points in time, and the conditions for mixed-strategy equilibrium are most likely to obtain.

In this early phase of the game, the tournament effects arising from only needing to score more runs than one's opponent to win (i.e., that winning by one run is as good as winning by ten) are minimized, and in such a case maximizing the expected number of runs scored in the inning is a good approximation to a team's true objective of maximizing the chances of winning the game. Palmer and Thorn [6] present a table of the expected number of future runs in an inning as a function of the current number of outs and location of runners. The expected number of runs with a runner on first base and no outs is 0.783 ($\approx v_B$). The expected number of runs with a runner on second base and no outs (i.e., the situation after a successful attempt to steal) is 1.068; the expected number of runs with no runners on and one out (i.e., the situation after a failed attempt) is 0.249. Therefore,

$$\frac{v_B - v_F}{v_S - v_F} \approx \frac{0.783 - 0.249}{1.068 - 0.249} = 0.652.$$

The theory thus predicts a relationship of approxmiately $\pi^\star = 0.652 + 0.348 p_S^\star$.

To test this prediction, play-by-play data from all Major League Baseball games played in the 1974 through 1992 seasons, inclusive, were examined to identify all situations in which the first hitter on a team reached first base in his team's first turn at bat.[2] For each player in each season, the number of times he reached first base in this situation, the number of times he attempted to steal, and the number of times he was successful were tabulated.

Three specifications of the relationship between the frequency of attempting to steal p_S and success percentage π are investigated: $\pi = \alpha_0$; $\pi = \beta_0 + \beta_1 p_S$; and $\pi = \delta_0 + \delta_1 p_S + \delta_2 p_S^2$. In practice, both p_S and π in this model are observed with randomness, as the decision whether to attempt a steal is being modeled as the product of an equilibrium in mixed strategies, and the realization of π depends on the realizations of the defense's mixed strategy. For the purposes of the estimation, the values of p_S are treated as being observed without error, while the observations of π are treated as being observed with noise. Since the noise in the observation of π decreases as the (square root of the) number of attempts made, and the noise in the oberation of p_S decreases as the (square root of the) number of opportunities, the noise in the observation of π is more salient. Furthermore, since the number of attempts made by different players is different, there is heteroskedasticity in the observations of π across different player-seasons. To accommodate this, a maximum likelihood approach is used.

2. These data are available on the website of Retrosheet, `http://www.retrosheet.org`. The term "play-by-play" means that the dataset identifies all events that change the number of outs, or the configuration of baserunners; that is to say, all batter outcomes, as well as the timing and outcome of stolen base plays.

A cutoff of a minimum of 20 observations for a player-season was used, resulting in 326 player-seasons; adjusting this cutoff does not substantively change the results. A scatter-plot of the data is shown in Figure 1. The first row in Table 2 presents the estimated relationships between π and p_S for the three specifications. Negative log-likelihoods are given in parentheses. The results indicate that the affine model significantly outperforms the constant specification, and the quadratic model is not significantly better than the affine model. The coefficients in the affine model are close to those predicted by the theory.

Many runners appear in this sample appear multiple times in different seasons during their career. As such, panel data concerns may arise. To address this, the same three specifications were estimated using career-level data for players. Again, a cutoff of a minimum of 20 observations for a player over his career was used, resulting in 299 players.[3] The scatterplot of this data appears in Figure 2. The second row of Table 2 presents the estimates for the models on this data. No qualitative differences in coefficient estimates are observed in comparing the player-season to player-career data, and, again, the affine model significantly outperforms the constant model, and the quadratic model can be rejected in favor of the affine model. The coefficient estimates for the affine model again track the theoretical prediction closely.

The parameter π is affected not only by the talents of a runner, but also by characteristics of a pitcher. Therefore, viewed from the defensive perspective, the same relationship should hold across the population of pitchers. There is an asymmetry in how pitchers are used in a baseball season compared to other regular players: while a regular player will appear in almost all of his team's games, pitchers are used in a rotation, pitching every fourth or fifth game on average. Therefore, the number of observations in the dataset for a given season for a pitcher is small, making individual season analysis impractical.[4] Aggregating over pitchers' careers, however, a similar pattern emerges to the results for runners. Again using a minimum of 20 observations, the scatterplot for the 309 pitchers appears as Figure 3. The third row of Table 2 presents the three models estimated using pitchers' career data. Again the coefficients are qualitatively similar to those of both estimations using runners' data, as well as the predictions of the theory.

3. This number is less than that for the number of player-seasons because many players appear multiple times in the season-by-season sample.

4. Very few pitchers even participate in 10 game situations that meet the requirements in a season.

	N	constant	predicted specification	quadratic
Runners (season)	326	.76 (515.80)	$.63 + .33p_S$ (493.89)	$.64 + .35p_S - .06p_S^2$ (493.54)
Runners (career)	299	.74 (456.66)	$.62 + .37p_S$ (419.87)	$.62 + .37p_S - .03p_S^2$ (419.19)
Pitchers (career)	309	.73 (551.20)	$.65 + .30p_S$ (542.71)	$.66 + .24p_S + .09p_S^2$ (542.63)

Table 2. Estimated relationships between p_S and π. Numbers in parentheses are negative log-likelihoods.

5 Evidence of the defense's strategy

An increase in the baserunner's stealing skills results in increased attention paid to the runner by the defense. Recall that the inequality $v_R > v_B$ is motivated by an assumption that batters perform better when the defense chooses to focus on the runner (to play strategy R) relative to when the defense chooses to focus on the batter (to play strategy B). Since an increase in ρ causes the defense to increase its equilibrium probability of choosing R, the observed performance of the batter should improve.

This indicates that batters' performance should improve more with a runner on first base who has a higher ρ. Similar to the choice of the receiving player in Walker and Wooders, the defensive choice of B or R is not directly observed. Because of the presence of the theorized effect on batter performance, this prediction provides a way of *indirectly* observing the effect on the defense's choices.

This prediction is tested by considering pairs of two players from the same team. In each pair, the performance of one player while batting with the other player occupying first base is tabulated. Additionally, the performance of the batting player is also tabulated for situations in which he batted with no runner on base. There were 722 such batter-runner pairs in the dataset in which the runner was on first base at least 50 times when the batter came to bat.

To operationalize what it means for a batter to perform better, the sum of a player's on-base percentage and slugging percentage is used to index performance. This sum, abbreviated OPS in baseball, correlates highly with run-scoring at a team level.[5] Figure 4 contains a scatterplot of the data, with the frequency of stolen base attempts p_S on the horizontal axis, and the increase in the batter's performance ΔOPS on the vertical axis.

Group	N	ΔOPS Median	ΔOPS Mean	$\Delta\text{OPS} > 0$
$p_S = 0.0$	42	-0.004	0.006	20 (47.6%)
$p_S \in (0.0, 0.2]$	595	0.042	0.051	368 (61.8%)
$p_S > 0.2$	85	0.134	0.109	59 (69.4%)

Table 3. The change in batter performance, measured by OPS, with a runner on first base compared to with the bases empty. The four groups are delineated by the frequency with which the runner on first base attempted to steal during the season. N is the number of batter-runner pairs in each group.

A regression of ΔOPS on p_S gives the line (standard errors in parentheses)

$$\Delta\text{OPS} = \underset{(0.010)}{0.032} + \underset{(0.081)}{0.249}\, p_S$$

with a standard error of 0.187 and an adjusted R^2 of 0.011. Both coefficients are significantly different from zero at the .01 level. The positive slope agrees with the prediction of the model that batters hitting with a runner with higher ρ enjoy better batting performance. In this data, the relatively large standard error, visible also in the scatterplot, arises because even a sample of 50 times batting is not very large for evaluating an individual batter's performance.

Another way of looking at the data is to group the batter-runner pairs according to the frequency with which the runner attempted the stolen base play overall. Table 3 presents the data aggregated into three groups: a group in which the runner never was observed to attempt the stolen base play ($p_S = 0$); a group where the runner attempted with a positive frequency, but no more than twenty percent ($p_S \in (0.0, 0.2]$); and a group where the runner attempted more than twenty percent of the time ($p_S > 0.2$). Among these groups, the median and mean ΔOPS increases as the frequency with which the runner attempted the stolen base play increases. The percentage of the pairs in which the batter's realized performance was better with the runner on first base than without increases as p_S increases. The test of equality of the proportions of increases between the $p_S = 0$ and $p_S \in (0, 0.2]$ groups gives a p-value of .165; the test of equality of the proportions of increases between the $p_S \in (0, 0.2]$ and $p_S > 0.2$ groups gives a p-value of 0.066.

5. For the 1974 through 1992 seasons, the correlation between a team's OPS in a season and its runs scored per game was 0.935.

Finally, note the group where $p_S = 0$. It can be said that, for these runners, the stolen base play was not a salient activity, and for them, the mixed-strategy theory presented here is not appropriate.[6] In that case, there is no significant observed difference, on average, between the performance of the batters with and without those runners on base.

6 Conclusions

The theory of mixed-strategy Nash equilibrium in a stylized inspection game organizes the observations of baseball teams' conduct in the stolen base play in Major League Baseball. Complementing the earlier papers of Walker and Wooders in tennis and Chiappori et al in soccer, the results indicate that this theory is a useful descriptive tool when experienced, motivated players interact in an environment where asymmetric information is not a significant factor.

The structure of the dataset allows investigation of the effects of individual heterogeneity in a way that was not available to Chiappori et al in the data on soccer penalty kicks. However, the data is not sufficient to test, as in Walker and Wooders, whether behavior is serially uncorrelated. As a descriptive theory, though, it is perhaps less important whether objective randomization is truly occuring, as interpreting the theory in terms of beliefs about the other side's behavior may be sufficient.

A key advantage to the tennis, soccer, and baseball datasets is that in modern professional sports, informational asymmetries can certainly be said to be small. The modern athlete (and team, in team sports) has the means to closely study the skills and behavior of opponents, thereby making the underlying assumption of common knowledge of the game compelling. An emerging feature of the baseball dataset, however, is that it is currently being researched back over the more than a century of continuous existence of professional baseball in the United States. This extends back before the days of television, radio, and heavy scouting, to a time when it was not uncommon for a team to come to town with some completely unknown players. A preliminary analysis of partial data from this era indicates that the model presented here fails to describe the data in this setting where common knowledge almost certainly fails.

6. In particular, for them, it is likely the assumption that the play is always successful when the defense chooses B is violated, and therefore the strategy S may be dominated.

Acknowledgments

The author wishes to thank Ray Battalio, Amy Glass, Dror Goldberg, Mike Nelson, Steve Puller, John Straub, and John Van Huyck, as well as seminar participants at Texas A&M University, the Stanford Institute for Theoretical Economics, the 2003 North American Summer Meeting of the Econometric Society at Northwestern University, and the 2004 Second World Congress of the Game Theory Society in Marseille for helpful comments and suggestions on various drafts of this paper. The raw play-by-play data used in this paper are made publicly available on the website of Retrosheet (`http://www.retrosheet.org`); the author thanks David Smith for access to some files prior to their public release. The processed data files can be obtained from the author's website (`http://econweb.tamu.edu/turocy`). All errors remain the responsibility of the author.

Bibliography

[1] R. Bellman. Dynamic programming and Markovian decision processes, with application to baseball. In S. P. Ladany and R. E. Machol, editors, *Optimal Strategies in Sports*, pages 77--85. Elsevier-North Holland, New York, 1977.

[2] P.-A. Chiappori, S. Levitt, and T. Groseclose. Testing mixed strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92:1138--1151, 2002.

[3] G. P. Lindsey. A scientific approach to strategy in baseball. In S. P. Ladany and R. E. Machol, editors, *Optimal Strategies in Sports*, pages 1--30. Elsevier-North Holland, New York, 1977.

[4] C. M. Mottley. The Application of Operations-Research Methods to Athletic Games. *Journal of the Operations Research Society of America*, 2:335--338, 1954.

[5] D. W. Smith. Maury Wills and the value of a stolen base. *Baseball Research Journal*, 1980.

[6] John Thorn and Pete Palmer. *The Hidden Game of Baseball*. Doubleday & Company, Inc., Garden City NY, 1984.

[7] M. Walker and J. Wooders. Minimax play at wimbledon. *American Economic Review*, 91:1521--1538, 2001.

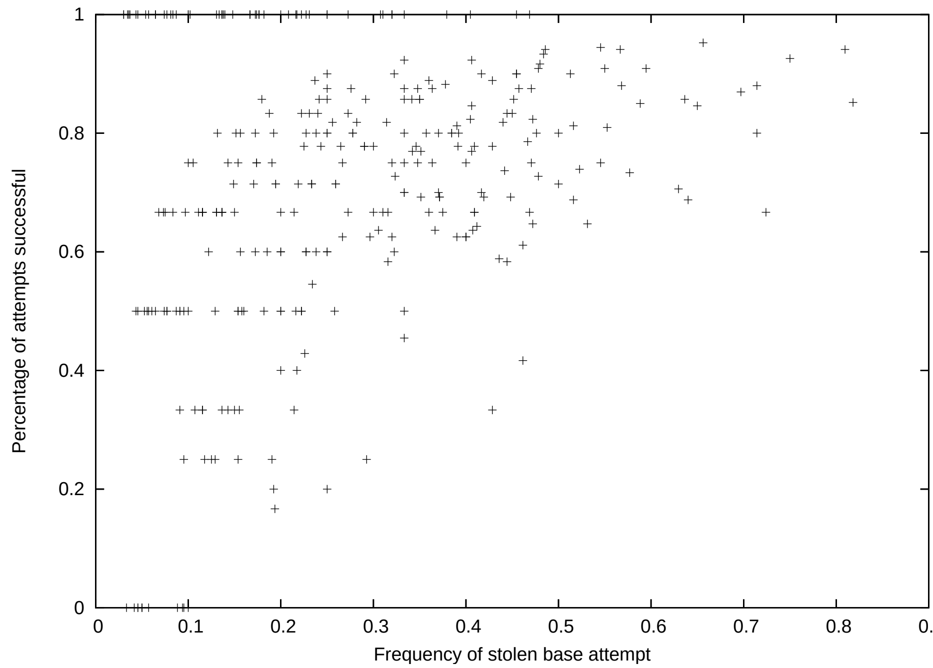

Figure 1. Scatterplot of frequency of attempt of stolen base play versus percentage of attempts successful. Each point represents one runner in one season, with a minimum requirement of having 20 opportunities (as defined by the game situation in the text).

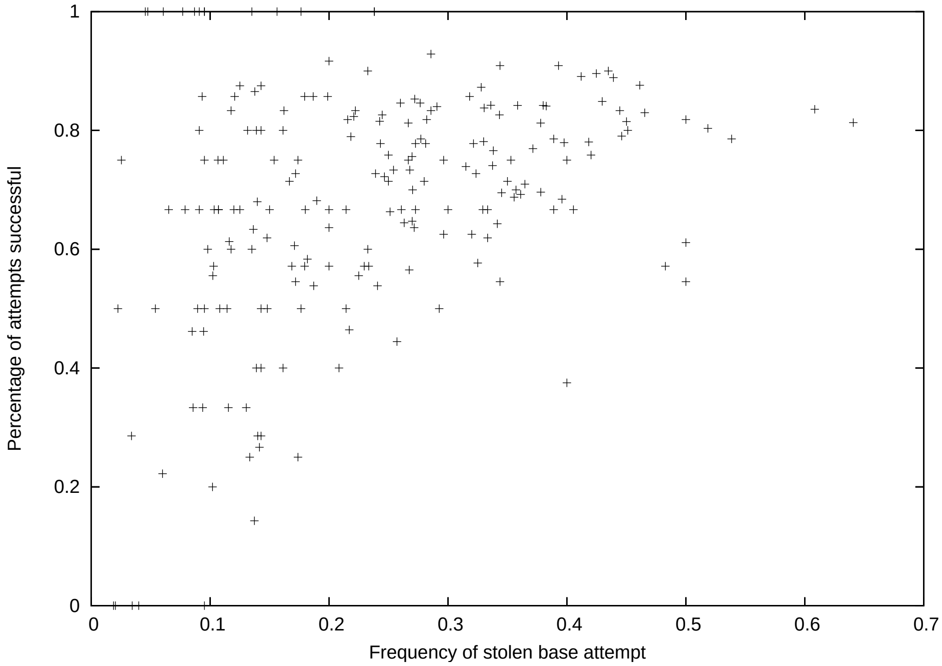

Figure 2. Scatterplot of frequency of attempt of stolen base play versus percentage of attempts successful. Each point represents one runner over his career, with a minimum requirement of having 20 opportunities (as defined by the game situation in the text).

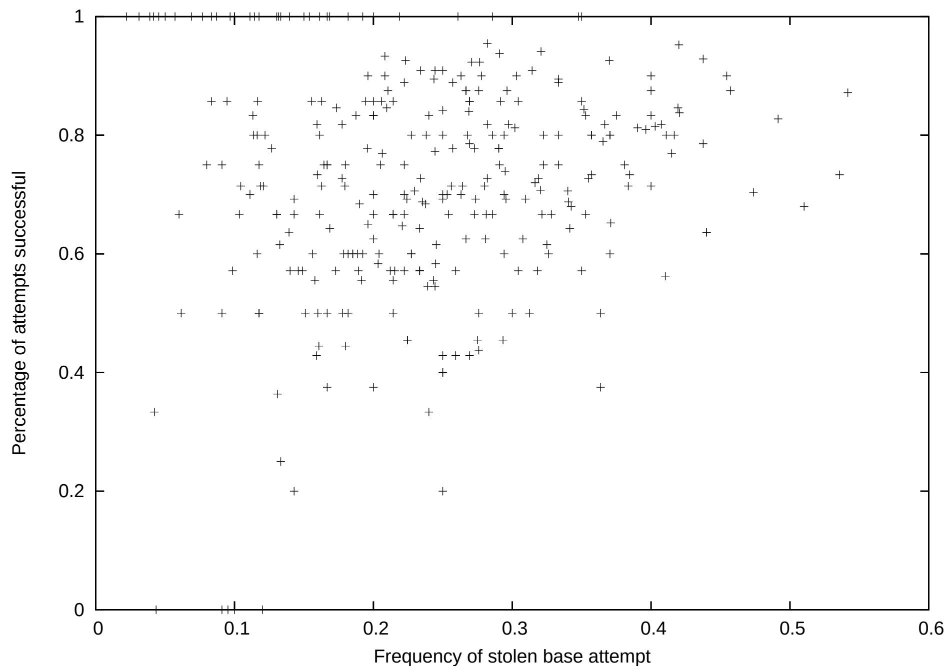

Figure 3. Scatterplot of frequency of attempt of stolen base play versus percentage of attempts successful. Each point represents one pitcher over his career, with a minimum requirement of having 20 opportunities (as defined by the game situation in the text).

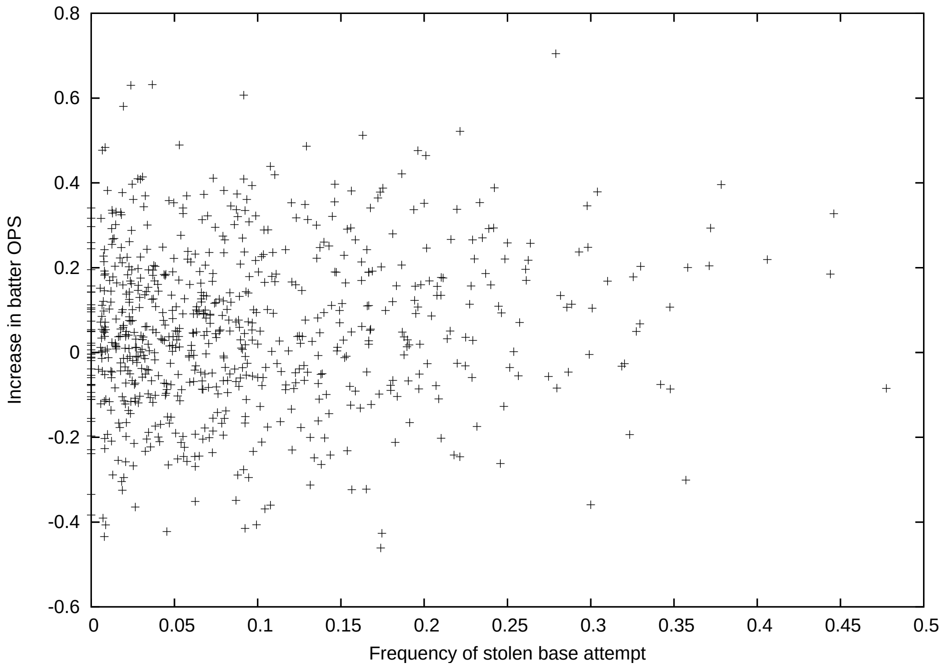

Figure 4. Scatterplot of frequency with which runner attempts to steal against the change in batter's performance with that runner on first (relative to batter's performance with no runner).