

A Winning Combination? Economic Theory Meets Sports

By Kristie M. Engemann and Michael T. Owyang



Unlike researchers in the natural sciences, economists often lack the ability to conduct laboratory or controlled experiments to test theories or make inference. In recent years, economists have begun to study “natural experiments”—naturally occurring events that provide a researcher with a basis to analyze outcomes within a clearly defined setting. One such artificial laboratory that economists have discovered is sports. Economists have used data from sports to examine such diverse issues as (1) risk behavior, (2) market efficiency, (3) market power and (4) discrimination.

Risk Behavior: Does Maximization Predict Coaches' Decisions?

A basic assumption in economic models is that, in competitive markets, firms maximize profits. In the sports world, such maximiza-

tion might be seen as a coach maximizing his team's chance of winning. Economist David Romer tested whether coaches make the optimal choice in a fourth-down situation in the National Football League (NFL). He argued that this analysis should be similar to the firm maximization model because winning is highly valued, coaches have pressure to win due to the competitive nature of the job and teams can learn from past experiences. Romer studied all of the regular-season NFL games during the 1998, 1999 and 2000

seasons, but used only the first quarter of the games in his analysis; later in the game, teams may change their strategy based on the score. Therefore, the first quarter should yield the best insight as to whether teams maximize their chances of winning. Romer analyzed the expected payoff from going for a first down on the fourth down at every point on the field versus kicking the ball (punting or a field goal attempt). His analysis compared the expected values of the outcome of a play, as well as the expected value of leaving the opponent with the ball at that spot on the field.

After taking all results into account, he estimated that teams are better off going for a first down than punting if they have fewer than four yards to go in their half of the field; if they have fewer than 6.5 yards to go on the other team's 45-yard line; and

if they have fewer than 9.8 yards to go on the other team's 33-yard line, at which point teams are within typical field-goal range. After the other team's 21-yard line, the value of going for it frequently outweighs the expected value of kicking a field goal, and at the 5-yard line, the team is always better off going for the first down or touchdown.

How did Romer's predictions compare with actual plays in the NFL games? In situations where teams were expected to be better off kicking the ball on fourth

down, they went for a first down less than 1 percent of the time. However, when teams were expected to be better off going for a first down, they kicked the ball almost 90 percent of the time. Romer estimated that if a team optimized in these situations throughout the whole game, it would win one more game every three seasons.

Romer surmised that coaches' previous experiences might cause more conservative decisions than one would predict using standard assumptions about optimizing behavior. Alternatively, a coach's objective might be more complicated than simply choosing plays that would result in the highest expected outcome. For instance, he might view activities that decrease the chance of winning (e.g., a failed first-down attempt) more negatively than he views a successful activity positively, which could stem from fan or owner preferences.

Risk Behavior: Does Game Theory Predict Player Behavior?

What happens when only two players are involved rather than entire teams? Economists Pierre-André Chiappori, Steven Levitt and Timothy Groseclose tested whether kickers and goalies used mixed strategies (i.e., chose strategies at random) to optimize their chances of being successful during penalty kicks in soccer. Even though soccer is a team sport, the penalty kicks involve just those two players and thus allow for a test of economic game theory. For each penalty kick, the kicker should maximize his chance of scoring, while the goalie should maximize his chance of preventing a score.

The authors studied all penalty kicks over a two-year period in the elite French league and over a three-year period in the elite Italian league. For each penalty kick, they had

the names of the kicker and the goalie, the direction the kicker kicked (right, left or center) and which foot he used, and the direction the goalie jumped (right, left or center). Due to the high speed of the ball, each player must decide which direction to kick/jump before the other player makes a move.

The authors contended that a goalie's strategy should depend on the kicker's past kicks, but a kicker's strategy should be independent of the goalie. In the authors' sample, when both players chose the kicker's natural side (which is the left side in most cases because the kickers usually kick with their right foot), the kicker scored 64 percent of the time, and when both chose the kicker's non-natural side, the kicker scored about 44 percent of the time. When the goalie jumped to the wrong side, the kicker scored 94 percent of the time when he kicked to his natural side and 89 percent of the time when he kicked to his non-natural side. Obvious from these data is that kickers are substantially more successful when kicking to their natural side.

The authors showed that, over their sample of 459 penalty kicks, the players used mixed strategies that one would expect in order to maximize behavior. Indeed, a kicker went to the center more often than the goalie did (17 percent versus 2 percent), and a kicker went to his natural side less often than the goalie did (45 percent versus 57 percent). Both players were more likely to go to the kicker's natural side than his non-natural side, and the case where they simultaneously went to the kicker's natural side was the most common (25 percent), followed by the goalie jumping to the kicker's natural side but the kicker going to the opposite (21 percent).

Based on all of these results, Chiappori, Levitt and Groseclose could not rule out that soccer players successfully optimize their behavior during penalty kicks.

Do Markets Work?

Many aspects of the sports labor market have been analyzed. The competitive environment of sports provides a setting in which one would expect merit-based outcomes to prevail.

Along this line, researchers Edward Fee, Charles Hadlock and Joshua Pierce studied promotions among coaches in the NFL.

The authors assessed whether promotions within teams were based on different criteria than promotions from outside. The researchers focused on promotions of offensive and defensive coordinators (level 2 coaches) to head coaches (level 1 coaches).

Examining data for NFL coaches from 1970 to 2001, the authors used a team's winning percentage as a measure of team performance. The authors used points scored for offensive coordinators and points allowed for defensive coordinators as a measure of individual performance. In assessing promotions of level 2 coaches to head coach on another team, the hiring decision depended on individual performance rather than team performance. In contrast, both team performance and individual performance mattered for promotion to head coach on the same team. A strong team performance actually decreased the likelihood of such a promotion, mostly because teams with winning records were less likely to replace their head coach. After controlling for the team, the highest individual performers were more likely to be promoted. However, the two effects essentially canceled each other out, leaving virtually no effect of individual performance on internal promotions.

Fee, Hadlock and Pierce likened this situation to top management at firms.¹ For senior management excluding CEOs, strong performers are more likely to obtain the position of CEO at a different firm rather than their own firm due to "slot constraints." The authors' findings do not support the theory that internal promotions serve as incentives for the best performers, at least not for top-level positions.

Market Power and the Labor Market

A sports league can be viewed as a monopsony—there is one buyer but many sellers of a product (players' services, in this case). A sports league, such as Major League Baseball (MLB), has market power because it can pay players less than their contribution to the team generates in revenue. However, the league cannot exercise as much market power for players who are eligible for salary arbitration or free agency.²



ANNUAL AND TOTAL COMPENSATION DIFFERENCES FOR MLB GROUPS

COMPARISON GROUP IS PLAYERS WITH LESS THAN THREE YEARS OF SERVICE				
	Nonpitchers		Pitchers	
	Whites	Nonwhites	Whites	Nonwhites
Annual compensation (% difference)				
Players with:				
3 or 4 years of service	36	43	36	28
5 years of service	48	44	28	53
Free agency	44	35	4	31
Total compensation (% difference)				
Players with:				
3 or 4 years of service	36	36	34	27
5 years of service	62	56	54	76
Free agency	68	59	51	66

SOURCE: Kahn (1993).

To study the effect of free agency and salary arbitration on salary and contract length, economist Lawrence Kahn looked at all players from 1987 to 1990. He obtained data on each player's salary, contract, total compensation, performance statistics and local demographics (e.g., population and per capita income of the team's metropolitan statistical area) and performed separate analyses on nonpitchers and pitchers. Compared to players with less than three years of service, free agents earned about five times more each year and players with five years of service earned between five and six times more each year during this time period.

The table shows the results of Kahn's analysis after controlling for performance statistics, years of experience, etc. The top panel shows that players with arbitration and free agency earned higher annual salaries than players with fewer than three years of service. A notable exception is white pitchers who were free agents—they earned roughly the same amount as those with fewer than three years of service. However, free agents were the only group with consistently longer contract length, which would affect total compensation (bottom panel) perhaps more than annual salary alone.

Kahn argued that his results are in line with free agents' willingness to accept a lower annual salary for the insurance that comes with longer contracts. He also argued that the significant effect of free agency on contract length shows that teams are willing to sign longer contracts only at

the possibility of losing a player, thus avoiding a "bidding war" with other teams.

Discrimination in Pay

Many economists have studied discrimination in the labor market, including the sports labor market. Researchers Kahn and Peter Sherer examined pay differentials between white and black players in the National Basketball Association (NBA) during the 1985-86 season. In 1985-86, about 75 percent of players were black, and, on average, black players earned almost 3 percent more than white players. In fact, the only three players during that season who earned more than \$2 million were black (Magic Johnson, Moses Malone and Kareem Abdul-Jabbar). Also, white players tended to play in cities with lower population, a higher white share of population and higher home-game attendance.

Kahn and Sherer reached a different conclusion regarding pay differentials after controlling for players' performance statistics (e.g., points, minutes per game and number of seasons played), team characteristics (winning percentage and home attendance) and some local demographics (total population and the percentage black, and per capita income). White players earned about 20 percent more than black players, all else equal, in the mid-1980s. In addition, the authors found that a white player with the same performance level as a black player would bring in a total of 8,000 to 13,000 more fans at home games, which they estimated was an extra \$80,000 to \$130,000 in revenue. The authors argued that their results reflect customer discrimination (rather than employer or co-worker discrimination) because fans appeared to be willing to pay a premium to watch white players.

Economist Barton Hamilton studied whether this "white premium" continued into the 1990s by examining salaries during the 1994-95 NBA season. For this season, the average black player earned about 17 percent more than the average white player, and nine of the 10 highest-paid players were black. Like Kahn and Sherer, Hamilton controlled for players' performance, team characteristics and local demographics to determine the true impact of race on a player's salary. Unlike

the previous authors, he found no significant pay differential between the average white and black player. However, among the stars and the superstars (i.e., those players in the top 25 percent and top 10 percent of the salary distribution, respectively), whites earned about 18 percent more than blacks. Because the stars are the most visible players on a team, Hamilton argued that this pay differential continued to reflect customer discrimination.

Discrimination in Play

Economists Joseph Price and Justin Wolfers performed a different evaluation of discrimination in the NBA. They estimated the amount of racial bias from referees when calling fouls on players of the opposite race, which can influence the on-court performance of the players. They examined every NBA game from the 1991-92 to the 2003-04 seasons and obtained statistics for each player and the race of the (randomly assigned) three referees for every game. The economists were able to compare the number of personal fouls a player received depending on the racial composition of the officiating crew.

About one-third of the referees during this time period were black, and black players accounted for over 80 percent of total minutes played. At first glance, the data showed that black players had more playing time and fewer fouls per 48 minutes played (the “foul rate”) than white players. Moreover, players had slightly lower foul rates when the officiating crew was of the same race, on average. A more in-depth analysis with controls for various characteristics (e.g., player position, height, weight, all-star status) showed that the foul rate for black players increased by about 4 percent when all three referees were white rather than black. As a consequence, playing time and points scored decreased slightly. Overall, the authors found that 62 percent of the black referees appeared to have a pro-black bias (by calling fewer fouls on black players), while 78 percent of the white referees appeared to have a pro-white bias (by calling more fouls on black players).

With these impacts on individual players, the authors tested the effect that this apparent referee bias had on the most important



outcome: winning. During the sample period, the margin of victory was one point in 4 percent of the games; thus, the seemingly small referee bias could have a large effect on the overall outcome. Indeed, Price and Wolfers argued that changing the racial composition of the referees to match that of the players on the team would lead to an increase in winning percentage for the team with more time played by black players from 48.6 percent to 50.5 percent.³

Not Just Fun and Games

Because economists do not generally have the opportunity to conduct laboratory experiments, sports provide an excellent opportunity to test theories, ranging from the existence of discrimination in the labor market to whether firms or individuals optimize their behavior to achieve a certain goal. With the high level of data availability that sports provides, undoubtedly there will be more natural experiments to analyze. 

Kristie Engemann is a research analyst, and Michael Owyang is an economist, both at the Federal Reserve Bank of St. Louis. For more on Owyang's work, see <http://research.stlouisfed.org/econ/owyang/index.html>.

ENDNOTES

- ¹ However, head coaches have a higher average turnover rate (22 percent) than CEOs (10 percent).
- ² In the MLB, players are not eligible for free agency, which allows them to negotiate a contract with multiple teams, until they have six years of major league service. Players with three, four or five years of major league service are eligible for salary arbitration. Under salary arbitration, the player and the team each submit a final offer, and an arbitrator must choose one of them. See Kahn (1993).
- ³ One conclusion that Price and Wolfers drew is that a potential bias by the referees for their own race exists. They argued that because NBA referees are heavily scrutinized after each of their games, it is most likely an unconscious bias.

REFERENCES

- Chiappori, Pierre-André; Levitt, Steven; and Groseclose, Timothy. “Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer.” *American Economic Review*, September 2002, Vol. 92, No. 4, pp. 1138-51.
- Fee, C. Edward; Hadlock, Charles J.; and Pierce, Joshua R. “Promotions in the Internal and External Labor Market: Evidence from Professional Football Coaching Careers.” *Journal of Business*, March 2006, Vol. 79, No. 2, pp. 821-50.
- Hamilton, Barton Hughes. “Racial Discrimination and Professional Basketball Salaries in the 1990s.” *Applied Economics*, March 1997, Vol. 29, No. 3, pp. 287-96.
- Kahn, Lawrence M. “Free Agency, Long-term Contracts and Compensation in Major League Baseball: Estimates from Panel Data.” *Review of Economics and Statistics*, February 1993, Vol. 75, No. 1, pp. 157-64.
- Kahn, Lawrence M.; and Sherer, Peter D. “Racial Differences in Professional Basketball Players’ Compensation.” *Journal of Labor Economics*, January 1988, Vol. 6, No. 1, pp. 40-61.
- Price, Joseph; and Wolfers, Justin. “Racial Discrimination among NBA Referees.” National Bureau of Economic Research Working Paper No. 13206, June 2007.
- Romer, David. “Do Firms Maximize? Evidence from Professional Football.” *Journal of Political Economy*, April 2006, Vol. 114, No. 2, pp. 340-65.