2008

Do Firms Have Short Memories? Evidence From Major League Baseball

Andrew Healy
Loyola Marymount University, ahealy@lmu.edu

Recommended Citation

Healy, Andrew. 2008. "Do Firms Have Short Memories? Evidence from Major League Baseball," Journal of Sports Economics 9(4): 407-424.

DO FIRMS HAVE SHORT MEMORIES?

EVIDENCE FROM MAJOR LEAGUE BASEBALL

Andrew Healy*

Loyola Marymount University

May 27, 2007

*Department of Economics, Loyola Marymount University, 1 LMU Drive, University Hall 4229, Los Angeles, CA 90045; e-mail: ahealy@lmu.edu

I started working on this project when I was a graduate student at the Massachusetts Institute of Technology, and I thank the National Science Foundation for the support of a graduate fellowship at that time. For helpful comments on earlier drafts of this paper, I thank Jay Bennett and John Healy. For excellent research assistance, I thank Stephen Salinas and Gena Gammie. Finally, I thank Sean Lahman for helping me to access the contract and performance data. All errors are my own.

DO FIRMS HAVE SHORT MEMORIES?

EVIDENCE FROM MAJOR LEAGUE BASEBALL

Abstract

When deciding what salary to offer an employee, a firm needs to predict that employee's future productivity. Classical theory predicts that firms will effectively use the available information to choose an appropriate salary offer. Evidence from baseball contracts indicates, however, that memory-based biases influence salary offers. Consistent with insights from psychology and behavioral economics, salaries are affected too much by recent performance compared to earlier performance. All organizations do not suffer equally from short memories. The teams that achieve the most with the money that they spend also use past performance data most effectively.

JEL codes: J31, L83

Keywords: compensation, baseball, memory, behavioral economics

A growing, and now large, body of evidence indicates that economic agents often do not conform to the classical rational model of behavior. In many cases, bounded rationality better describes behavior (Simon, 1957; Kahneman 2003). Rather than carefully considering all available information, agents often rely on rules of thumb to solve problems, particularly when those problems are complex (Tversky and Kahneman, 1974). One rule of thumb that may cause significant inefficiencies involves limited memory. Psychological research shows that people tend to primarily remember only salient events (Thompson, Reyes, and Bower, 1979). If economic agents forget or ignore less salient information, they will fail to act optimally when making decisions which could be improved by that information. This paper shows that these limits to memory lead firms to make systematic mistakes in their salary offers.

Since performance and pay are easily measured, the market for baseball players offers an ideal laboratory for testing theories about compensation (Sommers and Quinton, 1982; Kahn 1993), including theories relating to the effect of limited memory on salary offers. Also, since salaries are so high, the costs of any mistakes are amplified. The presence of inefficient behavior in the market for baseball players would likely extend to other environments where the stakes are lower. The data shows that teams are susceptible to a specific kind of inefficient behavior. Teams reward players for performing well in the immediate past, ignoring other evidence of a player's quality from his earlier performance. When choosing salary offers, teams have short memories.

This finding will likely not surprise baseball fans. Many anecdotes suggest that players often receive excessively lucrative contracts after one anomalous good season. For example, Adrian Beltre, a third baseman for the Los Angeles Dodgers, had an exceptional season in 2004. After the 2004 season, despite the absence of any similar success in the preceding years, the

Seattle Mariners offered Beltre a five-year, $64 million contract. The contract paid Beltre as if he would continue to perform at his 2004 level, but he has instead reverted to his pre-2004 form.[1] While other anecdotes also suggest that players are excessively rewarded for performing well in the final years of their contracts, testing the hypothesis that teams have short memories requires a comprehensive analysis of players' salaries.

In this paper, I analyze salary and performance data for all Major League Baseball hitters who signed free agent contracts from 1985-2004. The data shows that a player's performance this year is predicted about 20% more strongly by his performances from two and three years ago than by his performance from last year by itself. In contrast, a player's salary this year depends only half as much on his performances from two and three years ago as on his performance from last year. In other words, for determining future performance, there is more information in a player's earlier performance history than in his performance last year alone. Salaries, however, respond much more strongly to performance in that most recent year than to the earlier performance history.

Teams are not equally prone to undervaluing earlier performance relative to recent performance. Controlling for total payroll, the teams that win the most games use past performance data most effectively. Only the unsuccessful teams put significantly too much weight on recent performance relative to earlier performance. One plausible explanation for this result is that well-managed teams are less susceptible to making memory-based mistakes in their salary offers.

What could cause memory-based biases to affect baseball salaries? Previous research in the psychology of memory offers a compelling explanation. People often access the most salient memories when making decisions. Reacting primarily to salient memories in this way

corresponds to the *availability* heuristic (Tversky and Kahneman, 1973; Mullainathan, 2002). The most salient memory about a player who just had a remarkable season may be his recent outstanding play, while other relevant performance data fails to stand out as much. Availability could have caused the Seattle Mariners, for example, to believe that Adrian Beltre's lone exceptional season more accurately described the player's skills than his previous years of unexceptional play.[2] In general, this sort of behavior could explain why teams fail to take into account a player's performance in earlier years when making a salary offer.

In Section I, I describe the methods used to measure the determinants of a baseball player's performance and salary. In Section II, I discuss the data used in the analysis. In Section III, I present the main empirical results. Section IV concludes.

I. Estimation Strategy

As in previous research using baseball data, I focus on a player's contribution towards winning games, through which the player affects a team's revenues (Scully, 1974; Sommers and Quinton, 1982; Fort and Quirk, 1992). The previous research indicates that, to maximize revenues, a team primarily needs to focus on winning games. To the extent that other factors influence salary offers, the estimation strategy only requires that those other factors are not related to changes in a player's performance history over time.

In this paper, I will focus entirely on hitters. A variety of measures capture different aspects of a hitter's value to his team. The number of home runs that a player hits, a player's on-base percentage, and his slugging percentage are three such measures. Previous research suggests that a simple measure called OPS, the sum of on-base percentage and slugging percentage, is an accurate measure of a player's offensive value (Albert and Bennett, 2001). To

measure a player's value, OPS performs nearly as well as more complicated measures and better than any measure of similar complexity (Hakes and Sauer, 2006). In this paper, I use OPS to measure a player's worth. All results are largely the same if a different measure of player performance is used.

To formalize these ideas, consider a player *i* whose performance in year *t* is OPS_{it}. I model the player's performance this year as a weighted average of his performance in the previous *k* years, other observed variables x_{it}, and unobserved factors captured by u_{it}.

$$OPS_{it} = \sum_{s=t-k}^{t-1} \phi_s OPS_{is} + \varphi x_{it} + u_{it} \tag{1}$$

Equation (1) states that our prediction of a player's performance this year is based on his past performances, with higher weight likely assigned to more recent years. The error term, u_{it}, arises because previous performance does not perfectly predict current performance.

To model player salaries, I use the log of salary, since the salary data is heavily skewed. Take the log of a player's salary to be a function of the same variables that predict his performance.

$$\ln(Salary_{it}) = \sum_{s=t-k}^{t-1} \beta_s OPS_{is} + \alpha x_{it} + \varepsilon_{it} \tag{2}$$

The error term, ε_{it}, in the above regression comes from any aspect of player value that is not included in the model.

The regression results consider a variety of variables to be included in x_{it}, and the results are not sensitive to which variables are included. The variables that are considered for inclusion in x_{it} are a player's number of at-bats in previous seasons, a player's age, the square of a player's age, and team and year dummies. The at-bats variables are included to account for the fact that, conditional on a player's average success in a season as measured by OPS, a team may offer

different salaries to a player who has played more often in recent years. The age variables enter the regression to allow salaries and performance to have different age profiles. The year and team dummies account for the effect of any league-wide changes in hitters' statistics in a given year and the fact that a player's statistics depend in part on the quality of his teammates and the stadium in which he plays. As shown in Section 4, the results are robust to a variety of specifications.

Equations (1) and (2) suggest the test that I will conduct in this paper. If teams are optimally using previous performance information to predict players' future performances, the coefficients that relate previous performance to current performance should exhibit the same pattern over time as the coefficients that relate previous performance to current pay. Suppose we estimate that a 1 point increase in last year's OPS causes current OPS to be 0.6 points higher. Also suppose we estimate that an increase in OPS from two years ago of 1 point causes current OPS to be 0.3 points higher. Current performance is then twice as sensitive to last year's performance as to performance from two years ago. If management acts optimally, its salary offers will also be twice as sensitive to a player's performance in the previous year as to his performance from two years ago.

II. The Data

To test hypotheses relating to how effectively teams use past performance data to make salary offers, I utilize data on free agent signings in Major League Baseball from 1985-2004. The term free agent refers to a player whose contract has expired and who is free to negotiate with all teams except his original team. The competition for highly-valued free agents can be fierce.[3]

For the analysis, I focus on free agent signings that occur between October and April, when regular-season games are not taking place and the vast majority of free agent signings occur. Focusing on offseason signings allows for a clear comparison between a player's performance before and after the contract is signed. Of those signings, the regressions in this paper are based on all 500 observations for which players have at least 200 at-bats in each of the previous three seasons. By using this sample, I analyze the performance and salary data for all players who have played a significant amount of time in each of the three years preceding the signing of the free agent contract.

The data on players' salaries and past performances is obtained from the Baseball Archive.[4] The past performance data includes measures of every standard statistic for each player. I construct a player's OPS by adding his on-base percentage and slugging percentage. The salary information refers to a player's base salary for the given years. To determine free agent signings, I use the data on baseball transactions compiled by Retrosheet. This data contains information on all free agent signings that occur from 1985 to 2004.

The data shows that the average salary has increased from $500,000 in 1985 to $2.8 million in 2004. Baseball salaries are also highly skewed, and have become more skewed over time. The median baseball salaries in 1985 and 2004 were $410,000 and $870,000, respectively. To ensure that the few high salaries do not have excessive influence on the regression results, the salary regressions use the log of salary as the dependent variable.

III. Results

In this section, I estimate how a player's performance history affects current performance and salary. The data shows that past performance predicts current performance and salary in

strikingly different ways. In addition, the data shows that these differences and the inefficiencies they imply do not occur for the teams that are generally managed more effectively.

Testing for Short Memories

Table 1 presents the results from estimating the performance and salary equations. Columns (1) through (4) pertain to the regressions that use performance as the dependent variable, and columns (5) through (8) refer to the regressions that use salary as the dependent variable. If a model includes a particular independent variable, there will be a coefficient for that variable in the table. The table shows that the results are robust to different specifications. In the discussion below, I will focus on the results in columns 4 and 8, which consider the specifications that include three years of previous performance and the age variables.

As shown in column 4, a player's performance is predicted only slightly more effectively by last year's performance than by his performance from two years ago. The regression of current performance on the performance from the three previous years gives a coefficient of 0.343 for last year's performance and 0.263 for performance from two years ago. The difference between the two coefficients is insignificant ($p = 0.308$). On the other hand, an increase in last year's performance is about twice as effective at increasing that player's salary as is an increase in performance from two years ago. The coefficient for last year's performance is 3.345, while the coefficient for performance from two years ago is 1.623. The difference between the two coefficients is significant ($p = 0.001$).

While teams could do somewhat better by assigning more weight to performance data from two years ago, they are even more ineffective at using earlier performance data. Performance from three years ago predicts performance this year with a coefficient of 0.15, a

little less than half the coefficient for last year's performance. The hypothesis that performance from three years ago is an insignificant predictor of current performance can be rejected (p = 0.004). In contrast, the coefficient for performance from three years ago on salary is 0.215, about one-fifteenth as large as the effect of last year's performance. Moreover, the hypothesis that performance from three years ago has no effect on current salary cannot be rejected (p = 0.512).

Table 2 directly compares the relative effects that performances in different years have on current performance and on salary. For example, for current performance, the predictive power of performance from two years ago is about 0.767 times as large as performance from last year. For salary, the predictive power of performance from two years ago relative is about 0.485 times as large as performance from last year. The difference between these two ratios is noticeable, but not significant (p = 0.186). In predicting current performance, the effect of performance from three years ago is about 0.437 times the effect of performance from last year. The estimated effect of performance from three years ago on salary is only 0.064 times the effect of last year's performance. The difference between the two ratios is marginally significant (p = 0.056).

Taken together, the effects of all three earlier years on performance and salary indicate the significant gap between how teams actually use past performance data and how they would optimally use that information. For current performance, the predictive power of performance from two and three years ago is about 1.204 times larger than performance from last year. For salary, the predictive power of performance from two and three years ago is only 0.549 times the performance from last year. The difference between these two ratios is significant (p = 0.010). If the goal is to predict how well a player will perform, there is more combined information in

that player's performances from two and three years ago than in his performance from last year by itself. Nevertheless, teams put about twice as much weight on last year's performance than the combined weight they put on a player's performances from two and three years ago.

Successful and Unsuccessful Teams

In Figure 1, I plot average wins against average relative payroll for the thirty Major League Baseball teams from 1985-2004. To calculate relative payroll in a season, I divide each team's total payroll in a season by the mean payroll for all teams. Then I take the mean relative payroll for each team across seasons. In Figure 1, I also plot the line obtained by regressing average wins against average relative payroll. Teams that are above the line have won more games than their payrolls would predict. Teams below the line have won fewer games than their payrolls would predict.[5] In the discussion below, I refer to the teams that have achieved more wins than their payrolls would predict as successful teams. I refer to the teams that have achieved fewer wins than their payrolls would predict as unsuccessful teams.

By the standard of getting the most wins out of the salaries it has paid, Oakland has been the most successful team in baseball. Oakland averaged 86 wins per season from 1985 to 2004, even though its average payroll during that time predicts only 79 wins. On the other end, Tampa Bay was the least successful team, averaging only 65 wins from 1998 to 2004, ten fewer wins than the 75 wins that its average payroll predicts.[6] A complete description of the 15 successful teams and the 15 unsuccessful teams is in Figure 1.

This classification makes it possible to test for differences in how successful teams utilize past performance information compared to how unsuccessful teams use that information. To conduct these tests, I estimate regressions (1) and (2) separately for the successful and

unsuccessful teams. The results are shown in Table 3. In the discussion below, I focus on the results in columns 2, 4, 6, and 8 that consider the specifications that include the age variables.

Table 3 shows that successful teams do a better job of making their salary offers match up with how past performance data predicts players' present performances.[7] This result comes from two different sources. First, compared to unsuccessful teams, successful teams put higher relative weight on performances from two and three years ago when determining their salary offers. Successful (unsuccessful) teams put a weight of 2.958 (3.667) on last year's performance, a weight of 1.832 (1.399) on performance from two years ago, and a weight of 0.407 (0.126) on performance from three years ago.

Second, successful teams sign players whose current performance is better predicted by last year's performance. Successful (unsuccessful) teams sign players whose performances are predicted with a coefficient of 0.359 (0.297) for last year's performance, a coefficient of 0.285 (0.252) for performance from two years ago, and a coefficient of 0.103 (0.227) for performance from three years ago. The coefficient on performance from three years ago is insignificantly different from zero for successful teams ($p = 0.203$) and significantly different from zero for unsuccessful teams ($p = 0.001$). One possible explanation for these differences is that unsuccessful teams are more susceptible to overpaying players who have good statistics in the most recent season due to luck. To illustrate this idea, consider the example of Adrian Beltre, the player who signed a lucrative contract after one anomalous excellent season and then did not perform as well in the following season. If unsuccessful teams are particularly prone to signing players after one anomalous good season, then the players that these teams sign will have their performances better predicted by earlier years than players signed by successful teams.

Table 4 confirms that only the unsuccessful teams show significant memory-based biases in their salary offers. The successful teams treat past performance data similarly to how they optimally would. For example, in predicting current performance (salary), the effect of performance from three years ago is about 0.287 (0.138) as large as the effect of performance from last year. The difference between the two ratios is insignificant ($p = 0.638$). Taken together, performances from two and three years ago have an effect on present performance that is 1.081 times the effect of than last year's performance by itself. Similarly, performances from two and three years ago have 0.757 times the effect on salary as last year's performance by itself. The ratios for performance and salary are similar and the difference is insignificant ($p = 0.642$).

In contrast, the unsuccessful teams use past performance data to determine salaries in a significantly different way than that information predicts future performance. For these teams, performance from three years ago predicts current performance about 0.764 times as well as the effect of performance from last year. When choosing their salary offers, though, unsuccessful teams put only about 0.034 times as much weight on performance from three years ago as on last year's performance. The difference between these two ratios is marginally significant ($p = 0.052$). Altogether, for unsuccessful teams, performances from two and three years ago predict present performance about 1.613 times better than last year's performance does by itself. Still, for these teams, performances from two and three years ago predict salary only 0.411 times as well as last year's performance by itself. The difference between these ratios is significant ($p = 0.046$). Relative to last year's performance, a player's earlier performance history predicts future performance four times more effectively than it predicts that player's salary. No such difference is present for the successful teams.

IV. Conclusion

To best use the dollars that they spend on salaries, organizations need to predict how well players will perform in the future. The data shows that organizations make systematic mistakes in how they make these predictions. Teams infer too much from players' performances in the most recent season relative to performances from earlier years. The organizations that make these mistakes are the same ones that generally fail to spend their resources well. Organizations that are otherwise more successful also use past performance data more effectively.

The sizeable mistakes that unsuccessful teams make could arise from two different hypotheses about player behavior in the final year of a contract. First, it may be the case that players are rewarded for good luck in the last seasons of their contract, as previous research has demonstrated that CEOs are rewarded for luck (Bertrand and Mullainathan, 2002). In support of this hypothesis, Albert and Bennett (2001) show that a player's performance in any given year reflects a great deal of luck. On the other hand, there may be certain players who exert greater effort in the final years of their contracts and certain teams repeatedly fail to recognize this behavior. The tendency shown in this paper for teams to excessively reward performance at the end of a contract would amplify a player's incentive to try harder in his contract year. Future research could attempt to determine whether teams are rewarding players for luck in the final years of their contracts or for extra effort that they exerted in their contract years.

References

Albert, Jim and Jay Bennett, *Curve Ball: Baseball, Statistics and the Role of Chance in the Game*. New York: Springer (2001).

Bertrand, Marianne and Sendhil Mullainathan, "Are CEOs Rewarded for Luck? The Ones Without Principals Are," *Quarterly Journal of Economics*, 116:3 (2001), 901-32.

Camerer, Colin, "Does the Basketball Market Believe in the 'Hot Hand'?" *American Economic Review*, 79:5 (1989), 1257-61.

Gray, Philip and Stephen Gray, "Testing Market Efficiency: Evidence from the NFL Sports Betting Market," *Journal of Finance*, 52:4 (1997), 1725-37.

Hakes, Jahn K. and Raymond D. Sauer, "An Economic Evaluation of the *Moneyball* Hypothesis," *Journal of Economic Perspectives*, 20:3 (2006), 173-86.

Kahn, Lawrence, "Free Agency, Long-Term Contracts and Compensation in Major League Baseball: Estimates from Panel Data," *Review of Economics and Statistics*, 75:1 (1993), 157-64.

Kahneman, Daniel, "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review*. 93:5 (2003), 1449-1475

Kahneman, Daniel and Amos Tversky, "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology*, 4 (1973), 207-32.

Lahman, Sean, *The Baseball Archive*, http://www.baseball1.com (Data accessed in August 2006)

MacDonald, Don N. and Morgan O. Reynolds, "Are Baseball Players Paid Their Marginal Products?" *Managerial and Decision Economics,* 15:5 (1994), 443-57.

Mullainathan, Sendhil, "A Memory-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 117:3 (2002), 735-74.

Quirk, James and Rodney Fort, *Pay Dirt: The Business of Professional Team Sports*. Princeton, NJ: Princeton University Press (1992).

Retrosheet, *Transactions*, http://www.retrosheet.org (Data accessed in August 2006)

Scully, Gerald, "Pay and Performance in Major League Baseball," *American Economic Review*, 64:6 (1974), 915-30.

Simon, Herbert, "A Behavioral Model of Rational Choice," in *Models of Man*. New York: Wiley (1957).

Sommers, Paul and Noel Quinton, "Pay and Performance in Major League Baseball: The Case of the First Family of Free Agents," *Journal of Human Resources,* 17:3 (1982), 426-36.

Thompson, W.C., R.M. Reyes, and G.H. Bower, "Delayed Effects of Availability on Judgment," Stanford University manuscript (1979).

Tversky, Amos and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185:4157 (1974), 1124-31.

Notes

[1]Statisticians often use on-base plus slugging percentage (OPS) to measure a hitter's performance (Albert and Bennett, 2001). From 1999-2003, Beltre averaged a .756 OPS with a highest value of .835 in 2000. Then in 2004, his OPS jumped to 1.017. In the two seasons since, Beltre had a .716 and .793 OPS, respectively.

[2]A related phenomenon, the "hot hand", can also be understood by invoking the availability heuristic as it applies to memory. Camerer (1989), for example, finds that bettors incorrectly believe that a basketball team's recent success predicts its future success. Likewise, Gray and Gray (1997) find that odds for football games are skewed towards teams that have had recent success. They find that bettors' behavior "is consistent with the idea that the market overreacts to recent form, discounting the performance of the team over the season as a whole."

[3]Baseball players gained the right to free agency in 1973. Even though owners of baseball teams have been found guilty of colluding to keep players' salaries down as recently as the early 1990s, there is sufficient competition in the market for baseball players that experienced players have been receiving approximately marginal product wages going back to at least 1986 (MacDonald and Reynolds, 1994).

[4]I thank Sean Lahman of *The Baseball Archive* for help in accessing this data.

[5]The division of teams remains the same if a higher-order polynomial in relative payroll is used in the regression, rather than a simple linear regression.

[6]Tampa Bay entered Major League Baseball as an expansion team in 1998.

[7]In addition, the R^2 for the salary regressions are larger for the successful teams than for the unsuccessful teams, indicating that there is less unexplained variation in the salary offers made by successful teams.

Table 1: How past performance predicts present performance and salary

	Dependent variable: Performance in the present year				Dependent variable: Salary in the present year			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Performance last year	.388	.345	.385	.343	3.414	3.355	3.436	3.345
	(.042)	(.045)	(.042)	(.045)	(.32)	(.342)	(.313)	(.336)
Performance two years ago	.324	.272	.323	.263	1.603	1.541	1.708	1.623
	(.046)	(.051)	(.045)	(.051)	(.345)	(.385)	(.337)	(.38)
Performance three years ago		.144		.15		.253		.215
		(.052)		(.051)		(.39)		(.384)
At-bats last year	.000	-.003	-.002	-.007	1.347	1.235	1.288	1.163
	(.024)	(.024)	(.024)	(.024)	(.183)	(.184)	(.179)	(.181)
At-bats two years ago	.037	.036	.03	.03	1.089	1.059	1.21	1.105
	(.029)	(.031)	(.029)	(.03)	(.219)	(.231)	(.215)	(.228)
At-bats three years ago		-.01		-.002		.609		.71
		(.036)		(.035)		(.269)		(.266)
Age			-.019	-.013			.574	.441
			(.015)	(.017)			(.115)	(.131)
Age^2			.0002	.0001			-.0086	-.0068
			(.0002)	(.0003)			(.0017)	(.0019)
Number of observations	536	500	536	500	536	500	536	500
R^2	.464	.488	.483	.508	.631	.637	.65	.65

Notes:

a) Regression standard errors are in parentheses. Regressions include all players who had at least 150 at-bats in the present year and the three previous years.

b) A player's performance is measured by his OPS, the at-bats variables are expressed in thousands of at-bats, and the salary variable is the log of a player's salary.

c) Each regression includes year and team dummies.

Table 2: Last year's effect compared to earlier years

	Effect on performance	Effect on salary	*p*- value for test of equality
Relative effects of performance in different years			
$\frac{\text{Two years ago}}{\text{Last year}}$	0.767 (.211)	0.485 (.132)	$p = 0.186$
$\frac{\text{Three years ago}}{\text{Last year}}$	0.437 (.172)	0.064 (.114)	$p = 0.056$
$\frac{\text{Two years ago + Three years ago}}{\text{Last year}}$	1.204 (.286)	0.549 (.154)	$p = 0.010$

Notes:

a) Standard errors are in parentheses.

b) Boostrapping using the coefficients and their variance-covariance matrix from the regressions in columns 4 and 8 of Table 1 give the standard error estimates for the ratios.

Table 3: Past performance and salary for successful and unsuccessful teams

	Dependent variable: Performance in the present year				Dependent variable: Salary in the present year			
	Successful teams		Unsuccessful teams		Successful teams		Unsuccessful teams	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Performance last year	.377	.359	.294	.297	2.943	2.958	3.638	3.667
	(.07)	(.068)	(.061)	(.061)	(.469)	(.461)	(.495)	(.489)
Performance two years ago	.288	.285	.264	.252	1.759	1.832	1.399	1.381
	(.082)	(.08)	(.068)	(.068)	(.55)	(.54)	(.551)	(.547)
Performance three years ago	.085	.103	.226	.227	.501	.407	.146	.126
	(.083)	(.08)	(.07)	(.07)	(.553)	(.544)	(.571)	(.564)
At-bats last year	.042	.033	-.018	-.018	1.488	1.389	1.232	1.189
	(.046)	(.045)	(.029)	(.029)	(.311)	(.308)	(.237)	(.235)
At-bats two years ago	-.027	-.03	.087	.078	1.241	1.394	.789	.722
	(.046)	(.046)	(.042)	(.042)	(.311)	(.309)	(.344)	(.341)
At-bats three years ago	.076	.076	-.101	-.092	-.035	.077	1.309	1.415
	(.056)	(.055)	(.048)	(.048)	(.377)	(.372)	(.389)	(.386)
Age		-.031		-.006		.564		.27
		(.027)		(.023)		(.183)		(.189)
Age^2		.0004		.0000		-.0084		-.0044
		(.0004)		(.0003)		(.0027)		(.0028)
Number of observations	241	241	259	259	241	241	259	259
R^2	.5	.532	.521	.531	.694	.708	.631	.643

Notes:

a) Regression standard errors are in parentheses. Regressions include all players who had at least 150 at-bats in the present year and the three previous years.

b) A player's performance is measured by his OPS, the at-bats variables are expressed in thousands of at-bats, and the salary variable is the log of a player's salary.

c) Each regression includes year and team dummies.

d) Successful teams are the teams that have wins above what their payrolls from 1985-2004 would predict.

Table 4: Relative effects for different kinds of teams

	Effect on performance	Effect on salary	p- value for test of equality
Relative effects of performance in different years			
A. Successful teams			
$\frac{\text{Two years ago}}{\text{Last year}}$	0.794	0.619	$p =$ 0.770
	(.365)	(.268)	
$\frac{\text{Three years ago}}{\text{Last year}}$	0.287	0.138	$p =$ 0.638
	(.268)	(.209)	
$\frac{\text{Two years ago + Three years ago}}{\text{Last year}}$	1.081	0.757	$p =$ 0.642
	(.482)	(.323)	
B. Unsuccessful teams			
$\frac{\text{Two years ago}}{\text{Last year}}$	0.848	0.377	$p =$ 0.352
	(.347)	(.188)	
$\frac{\text{Three years ago}}{\text{Last year}}$	0.764	0.034	$p =$ 0.052
	(.362)	(.166)	
$\frac{\text{Two years ago + Three years ago}}{\text{Last year}}$	1.613	0.411	$p =$ 0.046
	(.566)	(.234)	

Notes:

a) Standard errors are in parentheses.

b) Boostrapping using the coefficients and their variance-covariance matrix from the regressions in columns 2, 4, 6, and 8 of Table 3 give the standard error estimates for the ratios.

Figure 1: Team payroll and wins

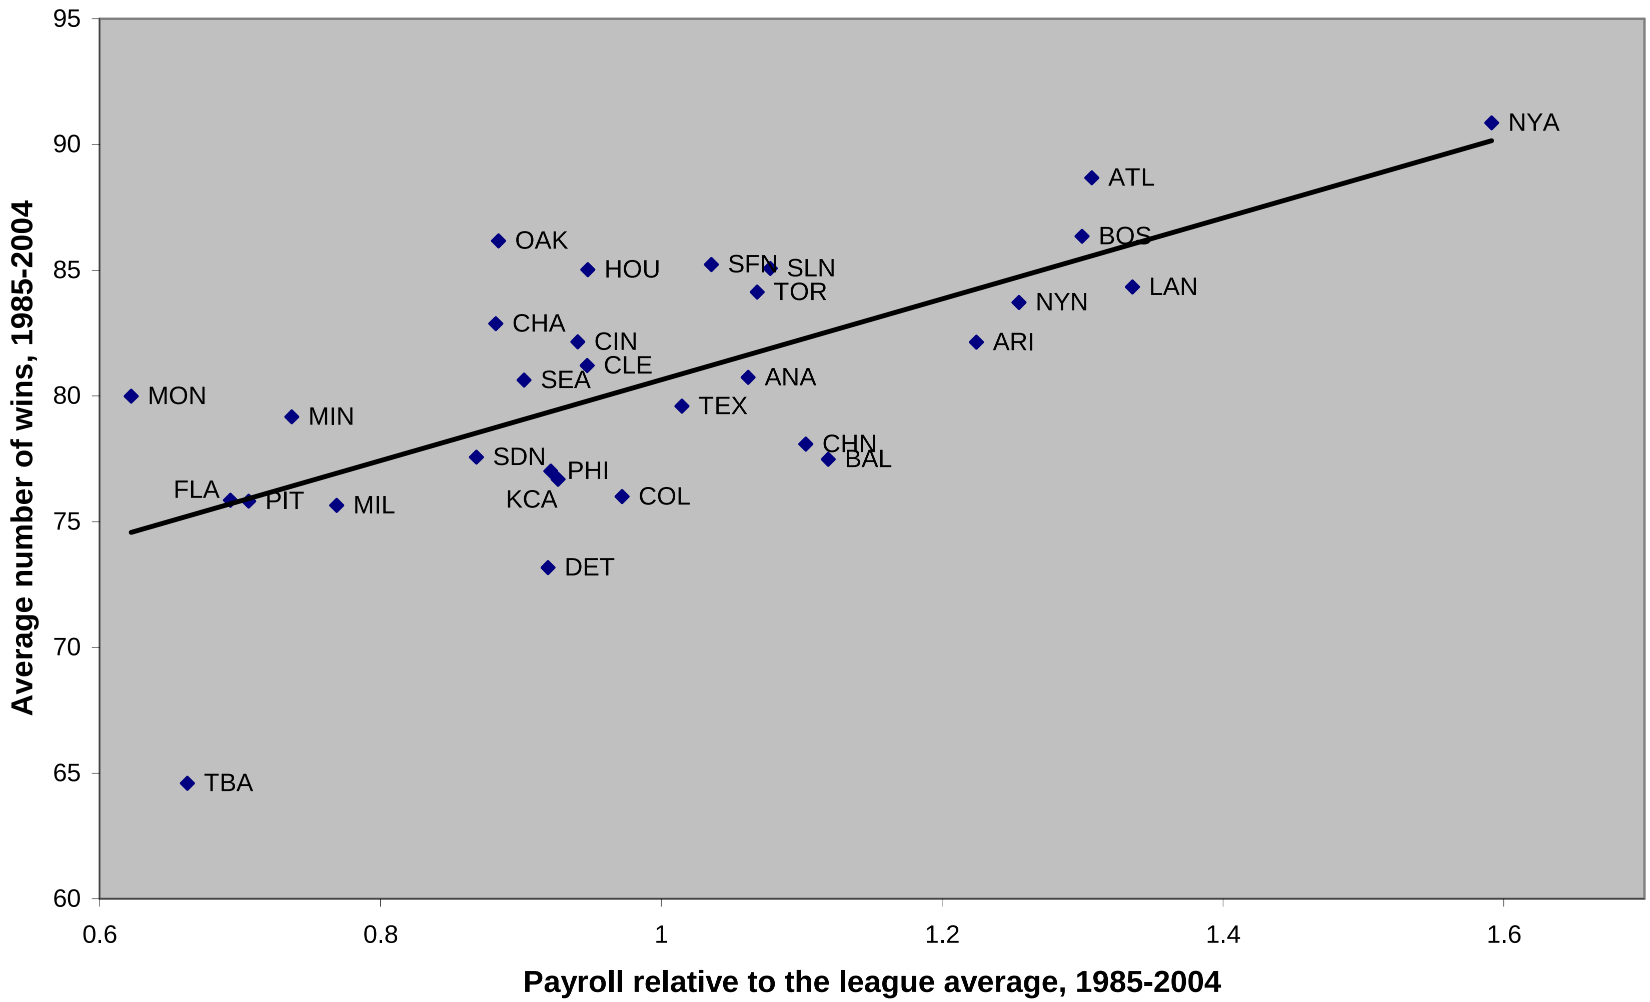

The teams above the line are Oakland (OAK), Montreal (MON), Houston (HOU), Chicago White Sox (CHA), San Francisco (SFN), St. Louis (SLN), Atlanta (ATL), Minnesota (MIN), Cincinnati (CIN), Toronto (TOR), Seattle (SEA), Cleveland (CLE), Boston (BOS), New York Yankees (NYA), and Florida (FLO).
The teams below the line are Pittsburgh (PIT), Anaheim (ANA), San Diego (SDN), New York Mets (NYN), Milwaukee (MIL), Texas (TEX), Los Angeles (LAN), Arizona (ARI), Philadelphia (PHI), Kansas City (KCA), Colorado (COL), Chicago Cubs (CHN), Baltimore (BAL), Detroit (DET), and Tampa Bay (TBA).