

Superstars and Journeymen: An Analysis of National Football Team's Allocation of the Salary Cap across Rosters, 2000-2005

Kevin G. Quinn[†], Melissa Geier^{††}, and Anne Berkovitz^{†††}

June 2007

Abstract

The National Football League constrains teams' payrolls via a "salary cap." We analyze how teams allocate cap spending across rosters using a data set of over 10,000 player-season observations during 2000-2005. We find that a few players account for relatively high portions of teams' caps, and that the players' "cap values" are consistent with both "superstar" and Yule-Simon income distributions. A theoretical model based on a utility function convex with respect to winning is used to explain this result. We also find that the cap has been substantially effective in reducing teams' ability to "spend their way to championships."

JEL Classification Codes: L83, J23, J42

Keywords: Sports, NFL

[†]St. Norbert College, Department of Economics, 100 Grant Street, De Pere, WI 54115, USA, kevin.quinn@snc.edu, phone: 920-403-3447, fax: 920-403-4098

^{††}St. Norbert College, Department of Economics, 100 Grant Street, De Pere, WI 54115

^{†††}St. Norbert College, Department of Economics, 100 Grant Street, De Pere, WI 54115

Introduction

The 1993 collective bargaining process between the NFL and its players union resulted in an agreement in which players earned some degree of free agency while owners got a salary cap, beginning with the 1994 season. The cap is actually a team payroll cap that requires that each team in the league keep the sum of player contract “cap values” below a maximum that varies from year to year. Because of the way in which the cap is accounted, an individual player’s cap charge against his team’s total cap may diverge significantly from his actual pay in any given season.

This paper reviews the history and the functioning of the NFL’s salary cap system, along with the rather sparse economics literature specifically related to NFL labor. A data set of 190 team-year observations that includes salary data for over 10,000 NFL player-seasons during six seasons (2000-2005) has been assembled for this paper. The data set is used to describe the manner in which teams allocate salary cap dollars across players on their roster. A formal theoretical specification of the general manager’s problem – maximization of the utility of winning subjected to the salary cap constraint - also is developed, and first order conditions are derived and analyzed. The implications of the model are tested against the data set, and some conclusions drawn about the functioning of the NFL player market. Finally, we consider relationships between winning and cap spending across teams and players.

The standard approach to modeling the economic decisions facing a firm is a constrained optimization problem in which output is chosen so as to maximize profit subject to technology and prices for both inputs and outputs. It is customary in the North

American sports economics literature to consider winning percentage as a team's output, and labor as its input. However, the existence of the salary cap system, along with the significant degree of revenue-sharing in the NFL, results in very little sensitivity of revenues or profits with respect to winning. On the other hand, it is apparent that most NFL team general managers do seek to win – in fact, the continuing employment of a general manager frequently appears to be strongly related to his or her team's competitive success. Therefore, a more realistic description of the annual constrained optimization problem facing NFL team managers may be the maximization of a utility function that depends on winning percentage; this optimization is performed subject to the annual salary cap. This is the premise of the theoretical model developed in this paper.

Literature Review

The question of whether or not payroll imbalance in a sports league has negative consequences for competitive balance is of considerable to sports economists. Hall, Szymanski, and Zimbalist (2002) consider the relationship between team payroll and competitive success in both Major League Baseball and English soccer, finding that higher payrolls result in more winning in soccer, but perhaps not in baseball. Larsen, Fenn, and Spenner (2006) determine that introduction of the salary cap in the NFL reduced the concentration of talent among a few teams, therefore improving competitive balance across the league.

Kowalewski and Leeds (1999) also study the NFL salary cap, finding that it led to increased inequality of pay among NFL players - higher pay for superstars at the expense

of rookies and marginal players. Leeds and Kowalewski (2001) use an isoquant-isocost output expansion path model in the context of a Nash bargaining solution to consider the variations in the returns to increasing skill levels of different players. They suggest that the nature of the NFL production function can result differing optimal pay levels for different skills as constraints on teams' payrolls change as a result of the salary cap.

Porter and Scully (1996) suggest that distribution of earnings among athletes may follow a rank-order tournament mode (Lazear and Rosen, 1981). They find that a variety of factors affect the distribution of athlete earnings, and that Gini coefficients vary between 0.22 and 0.64 among various team and individual sports. More generally, however, Rosen (1981) notes that in some labor markets, the revenue generation function is convex with respect to the quality of the output. When this is the case, those producers with relatively less productivity (talent) will find their output to be relatively poor substitutes for those with greater productivity, even if the talent differences are rather small. As a result, those with top talent, even if only slightly more so than those in the second tier, will earn significantly more; "the income distribution is stretched out in its right-hand tail compared to the distribution of talent" (p. 846). Rosen calls this the "superstar effect."

Chung and Cox (1994) claim that a particular class of distributions examined by Yule (1924) and Simon (1955) - special cases of Pareto's (1897) income distribution density function- fit the "superstar" distributions of artistic output. Most notably, they find that superstar income outcomes do not necessarily rely upon talent differences; rather, they find Yule-Simon income distributions even among individuals with equal talent. Cox and Hill (2005) apply this methodology to the distribution of the most highly-

paid NBA players for the 1997-1998 and 2000-2001 seasons, also noting that Scully's (1974) and Kahn and Sherer's (1988) use of log-linear specifications relating pay to talent are consistent with the philosophy and character of Yule-Simon distributions.

Finally, there is also some discussion in the sports economics literature about the nature of team owners' and managers' objective functions. Sloane (1971) and Kesenne (1996) examine the implications of win-maximization, while Quirk and Fort (1992) assume profit maximization.

NFL Capology and Revenue Sharing

The NFL's "hard" salary cap often is given substantial credit for the league's remarkable profitability compared to its peer North American leagues. The cap, which took effect with the 1994 season, was the result of the 1993 Collective Bargaining Agreement between the league and its players union. The cap is based on a guarantee percentage of the league's Defined Gross Revenues (DGR). During the period considered in this study (2000-2005), this percentage was about 56%.¹

The DGR essentially are the revenues that are shared evenly among teams, consisting of media revenues, licensing revenues, and the portion (about 40%) of gate revenues that are put into a pool and shared. Luxury seating and other stadium revenues were not included in DGR through the 2006 season. Figure 1 shows the portion of total team revenues represented by shared league revenues for the 32 NFL teams for the 2005 season. Shared revenues for 2005 were \$114 million per team, representing a mean of 59% of team's total revenues, and accounting for between 52% and 69% of total team

revenues for 24 of the 32 teams (Forbes, 2006; Spofford, 2006). The league revenue sharing system means that differences among NFL teams' operating profits depend far less on competitive success than on specific stadium arrangements with host cities.

Each year, the league-wide salary cap is determined by the league's estimated DGR multiplied by the players' guaranteed percentage. Each individual team's salary cap is derived by dividing this value by the number of teams in the league. Figure 2 indicates the salary cap per team by year since the inception of the cap.

A team may not enter into a set of player contracts that have cumulative annual cap values that exceed the team's annual limit. There also is a requirement that a team may not have a set of contracts with cap values that sum to less than 80% of the team's annual limit. The operative essence of the cap is that the sum of the "cap values" for a team's 53-man roster (45 of whom are considered "active" on any given game day, while 8 are considered "inactive") may not exceed the team's cap for that season, although there are some minor exceptions.² However, because of the nature of NFL contracts, the notion of player "cap values" is somewhat byzantine.

Player contracts can be for one season or for multiple seasons, and typically consist of a base salary as well as a variety of bonuses. Base salaries are not guaranteed - players may be cut at any time, with no further payments. Bonuses are lump sum payments that typically are not recoverable by teams (although this is beginning to change), and include signing bonuses (paid upon contract initiation), training camp bonuses (paid upon arrival at camp), roster bonuses (paid upon making the final 53-man roster just prior to the beginning of the season), and various performance bonuses as specified by the Collective Bargaining Agreement. Performance bonuses generally are

considered either "likely earned" (the player has met the performance threshold some time in his career) or "not likely earned" (the player has never met the threshold).

Going into each season, a player's annual "cap value" is given by the sum of the following: base salary, total contract signing bonus divided by the number of seasons of contract duration, training camp bonus, roster bonus, and "likely earned" performance bonus.

Cap values can vary among teams within a season. A player's "not likely earned" performance bonuses are not included in a player's cap cost, and therefore do not count toward a team's cap limit during the year earned. In the event that the player meets the performance threshold and is paid such a bonus, the team's cap for the following season is reduced by the resultant overage. In the event that a player is cut during a season (each season is treated as beginning on June 1), then the team's cap for the *following* season is reduced by the unamortized portion of the cut player's signing bonus. This is a common occurrence - team general managers appear to exhibit very high time discount rates, typically espousing a "win this year, and worry next year about next year" ethic. Figure 3 shows the relationship between mean team actual salaries paid in a given year to players and team cap charges that resulted from those payments. While the values are not equal, they are fairly close.

Salary Cap Data and Analysis

We have assembled a data set of 190 team-year observations for the 2000-2005 NFL seasons - 31 teams during 2001-02 and 32 during 2003-05 following the addition of the Houston expansion franchise. The data set includes only the top 55 players on each

team as ranked by salary, for a total of 10,436 player-year observations.³ The data were primarily drawn from USA Today and Fort (2006).

Table 1 shows the mean team winning percentages during the 2000-2005 period, along with teams' mean ratio of total cap spending to the league averages in each year. The maximum and minimum of these ratios during the 2000-2005 seasons is also shown by team. All but two of the means are between 94% and 106% of the league average. Individual season ratios range between 69% and 126%, with 28 of the 32 teams' maximum values less than 115% of that season's average, and 28 of 32 minimums of 75% or more. Table 1 also indicates the mean Gini coefficients of the distribution of player cap charges by team. The grand mean of these Ginis is 0.273, on the lower end of the range of those examined by Porter and Scully (1996).

Figure 4 shows the mean, maximum, and minimum individual player shares of team salary cap spending, ordered by rank on team (1st through 55th in terms of highest to lowest cap charge on team) for the 190 team-seasons represented in the data set. Figure 5 depicts the Lorenz curve associated with these data. Figure 6 shows the means of player cap shares of mean team cap spending by player cap rank across the data set against both a Pareto model and a ln-ln model. The natural log model provides an excellent description of these data for all but the top two highest cap cost players.⁴

Figures 7 and 8 show how teams chose to spend their cap money across different playing positions. While some positions represent on average greater shares of team cap spending, the distributions of each position by cap rank on team are qualitatively similar.

Theoretical Cap Allocation Model

We now turn to a theoretical construction of how team managers choose the allocation of fixed salary cap dollars across the N players on the roster. We will consider first a blended model wherein managers choose the roster (\mathbf{q}) so as to maximize utility, which in turn is a function of winning percentage (W) and profit (π):

$$\begin{aligned} \text{Max } U &= U[W(\mathbf{q}), \pi(W(\mathbf{q}))] \\ &\mathbf{q} \\ \text{s.t. } \quad &\Sigma c_i \leq C \text{ and } W = \Sigma q_i; \quad (i = 1, \dots, N) \end{aligned}$$

where c_i is the cap charge associated with player i on the roster. Note that this approach adopts the Scully (1974) assumption that team wins are the linear combination of contributions of individual players' talents; i.e., that there is no complementarity among players in the production of team wins.⁵

If we make the simplifying assumption that there is no relationship between winning and team profit – not a terribly unreasonable one for NFL teams during the period of interest here - then we can collapse our problem. This simpler problem is the manner in which to allocate a fixed pool of cap dollars across a specified number of players so as to maximize the manager's utility:

$$\text{Max}_{\mathbf{c}} U(W(\mathbf{q}(\mathbf{c}))) = U(\sum q_i(c_i))$$

$$\text{st: } W = \sum q_i + \varepsilon \text{ and } \sum c_i \leq C \text{ (} i=1, \dots, N \text{)}$$

We shall assume that utility increases at an increasing rate with respect to winning; i.e., $U_W, U_{WW} > 0$, and that the win production function increases at a flat or decreasing rate with respect to roster quality; i.e., $W_q > 0$ and $W_{qq} \leq 0$. The salary cap cost of players is assumed to increase at a flat or increasing rate with respect to player quality; i.e., $q_c > 0$ and $q_{cc} \leq 0$.⁶

We specify a Lagrangian,

$$\mathcal{L} = U(\sum q_i(c_i)) - \lambda[\sum c_i - C]$$

where λ represents the shadow price associated with the marginal relaxation of the team salary cap. A positive value of λ , associated with an interior solution (that the manager would not choose to spend its entire salary cap on a single player), implies that the salary cap indeed constrains manager choice. Assuming such an interior solution, we obtain the following first-order conditions:

$$\text{FOC (a):} \quad \partial \mathcal{L} / \partial c_i = 0 \Rightarrow (\partial U / \partial W)(\partial W / \partial q_i)(\partial q_i / \partial c_i) - \lambda = 0 \text{ for each } i$$

$$\text{FOC (b):} \quad \partial \mathcal{L} / \partial \lambda = 0 \Rightarrow \Sigma c_i - C = 0$$

FOC (b) is the expected result that the manager will choose to spend all of the team's cap allotment, but FOC (a) has somewhat more interesting implications.

Recalling that the win production function is modeled here to be a linear combination of talent on the roster, we note that $\partial W / \partial q_i = 1$. This implies that $(\partial q_i / \partial c_i) = \partial W / \partial c_i = 1 / MC_W$ (where MC_W is the marginal cap cost of winning), so we can rewrite FOC (a) as

$$\begin{aligned} (\partial U / \partial W)(\partial W / \partial q_i)(\partial q_i / \partial c_i) - \lambda &= 0 \\ \Rightarrow MU_W / MC_W &= \lambda \text{ for each player } i \end{aligned}$$

where MU_W is the manager's marginal utility of winning. Because we have assumed that U_{WW} is greater than zero – manager utility increases at an increasing rate with winning – the marginal cap cost of winning must also increase proportionally with increasing talent level; that is, when players on a team are ordered from best to worst, their cap charges fall off more quickly than linearly with respect to talent fall-off. As long as there are at least some managers in the league for which marginal utility of winning is an increasing

function, then player cap charges across the league will decrease with decreasing talent, and will do so at more quickly than linearly.

In fact, this is what we observe empirically in the NFL. Even if the talent level diminishes nearly linearly (or even shallowly, as suggested by the data on QB and RB productivity shown in Figures 9 and 10), we find that player cap costs fall off logarithmically (Figures 11 and 12). This phenomenon holds for players on individual team rosters as well as across the entire league.

In a sense, NFL player pay may be better explained with a “marginal win utility product” model than with the standard marginal revenue product model. If managers seek to win, and that first-place finishes and championships are highly-prized outcomes, then the exponential decline in cap costs across rosters does not require exponential drop-offs in talent level, but rather is a manifestation of the superstar effect. This is good news, because there appears to be little empirical support for the marginal revenue explanation. Player talent does not appear to drop off exponentially players when ordered from best to worst, nor do revenues seem particularly connected with winning percentage. Furthermore, we need not resort to a tournament model, a Leeds-Kowalewski output expansion graph model, or a winner’s curse explanation to account for steeply higher pay for slightly more productive players.

Winning and Salary Cap Spending Across Players

An implication of the model developed above is that if the salary cap “works” i.e., constrains a manager’s selection of talent, then there should be a positive shadow price associated with higher cap spending. Furthermore, a functional cap means that there

should be some regression to the mean in terms of cap spending – teams that spend more this year should be spending less in future years, and vice versa. We have the opportunity to test for this in our data set, as some teams exhibited actual cap spending above and below the league average for any particular season. Another question of interest is whether or not team differences in the distribution of cap spending across players is associated with differing levels of competitive success.

Table 2 shows the correlations between the teams' cap spending vis-à-vis the league mean for the season for the 2000 through 2005 seasons. Teams' cap spending in the 2004 and 2005 seasons generally are negatively correlated with spending in the 2000 through 2002 seasons, suggesting that there is some reversion to the mean, and that the cap does function as advertised, albeit imperfectly. Figure 13 shows teams' winning percentage versus their cap spending for the year; there appears to be a relationship between cap spending and winning.

Figures 13 and 14 show the relationship between teams' on-field success and the distribution of total team cap spending across players on their rosters. Winning teams appear to spend on average a bit more of their salary cap on the players ranked 15th through 30th in terms of cap spending, and less on players ranked 35th and higher, than do losing teams. However, specific description of this relationship remains elusive.⁷

Figure 15 shows the relationship between teams' cap spending (relative to the league average for the season) and the Gini coefficient of cap spending across players. A simple OLS regression estimate (not shown here) of the data in Figure 15 indicates that the coefficient on cap spending is positive and significant to the 1% level, but that the R^2 of the regression is only 7.3%.

The OLS regression models shown in Table 3 indicate that there indeed is some (but not much) competitive advantage gained when teams' cap spending exceeds the average in a given season. The coefficient on cap spending indicates that a 10% increase above the mean cap spending in the league corresponds to about half a game more won during the season, and the R^2 of the regressions is below 8%. Furthermore, teams' Gini coefficients of spending across players is not found to significantly explain differences in their competitive success.

The conclusion suggested by this analysis is that the NFL salary cap regime that was in place during 2000-2005 was generally successful in removing any substantial competitive advantage to teams associated with having higher roster spending.

Concluding Remarks

This paper seeks to analyze some of the economic and competitive outcomes associated with the functioning of the NFL salary cap during the 2000-2005 period; specifically, we have concerned ourselves here with teams' decisions about how to allocate salary cap dollars across players on their rosters. We find that there are significant differences between teams' spending on "superstars" versus their spending on journeyman players. In particular, we find that cap spending across players is consistent with other types of income distribution, in concert with the sense of both Pareto and of superstar models.

Given the very high proportion of teams' revenues that are independent of on-field performance, economists might expect much more free-riding than appears to have been observed among NFL teams. This begs the question of why hasn't the league

collapsed into the typical economic pathology of a collective farm? Perhaps the answer is suggested by the fact that the league bans corporate ownership of teams.

While NFL revenues are significant – on the order of \$10 billion annually – they are small compared to those enjoyed by the large conglomerates that have sought team ownership in other North American leagues. It is easy to imagine that an NFL team would simply be treated as another corporate business unit, and operated so as to maximize profit. On the other hand, it may be that individual ownership of NFL teams reinforces an “ego” element of winning.

We offer here a theoretical explanation of these observations that depends more on the desire to win than on profit-seeking. We find that if the utility associated with winning is convex – that is, if team managers’ utility increases at an increasing rate with respect to winning – then small differences in talent among players can lead to large differences in the portion of team cap spending on them. If owners’ egos indeed are wrapped up with their teams’ competitive success, we would predict that distribution of salary cap spending across players would look as it does.

Furthermore, the league’s salary cap appears to successfully reduce the ability of these owners to “spend their way to championships.” Moreover, while there may be some rather difficult-to-detect strategies in cap allocation across players to enhance winning, teasing them out of the available data remains elusive.




Figure 1: Individual NFL Team Total Revenues from Shared League Revenues, 2005 (derived from Forbes, 2006; Spofford, 2006).



Figure 2: NFL Team Salary Cap Values, 1994-2007 (nominal dollars).




Figure 3: Mean Actual Player Salary Payments vs. Team Cap Values, 2000-2005.

Team	Mean Win Pct	Mean Gini for Team	Mean Annual Ratio of Team Cap Spending to NFL Avg	Max Cap Spending Ratio	Min Cap Spending Ratio
PHI	0.677	0.245	1.13	1.25	1.01
PIT	0.672	0.264	1.05	1.14	0.95
IND	0.667	0.297	1.01	1.07	0.90
NE	0.654	0.290	1.00	1.05	0.88
DEN	0.635	0.302	0.96	1.06	0.87
STL	0.594	0.280	1.04	1.11	0.95
GB	0.594	0.237	0.99	1.06	0.89
NYJ	0.567	0.305	1.04	1.11	0.95
BAL	0.563	0.274	1.01	1.08	0.79
SEA	0.563	0.288	1.04	1.15	0.90
TB	0.563	0.295	1.03	1.11	0.91
MIA	0.563	0.242	1.00	1.04	0.97
TEN	0.542	0.291	0.98	1.05	0.71
KC	0.518	0.289	0.96	1.13	0.79
MIN	0.506	0.270	1.01	1.26	0.70
OAK	0.479	0.281	0.99	1.12	0.88
CHI	0.469	0.236	1.04	1.07	1.00
NO	0.469	0.278	1.01	1.12	0.82
JAC	0.469	0.309	0.97	1.06	0.83
ATL	0.464	0.278	1.01	1.08	0.94
CAR	0.459	0.268	0.96	1.05	0.89
WAS	0.457	0.290	0.92	0.99	0.85
NYG	0.448	0.261	1.00	1.10	0.91
SF	0.427	0.252	0.86	1.07	0.69
DAL	0.417	0.279	0.95	1.08	0.73
CIN	0.407	0.266	1.04	1.12	0.96
BUF	0.407	0.280	0.96	1.08	0.82
SD	0.406	0.257	0.94	1.06	0.85
CLE	0.354	0.252	1.00	1.14	0.84
ARI	0.313	0.256	1.02	1.12	0.92
DET	0.313	0.255	1.02	1.18	0.89
HOU	0.281	0.270	1.06	1.18	0.96

Note: Means, maxes and mins refer to individual teams among the 2000-2005 seasons.

Table 1: Comparison of Win Percents, Gini Coefficients, and Cap Spending Among Teams, 2000-2005 Seasons;




Figure 4: Mean, Maximum, and Minimum Player Cap Shares of Team Totals, 190 NFL Team Seasons, 2000-2005




Figure 5: Lorenz Curve, Player Cap Shares of Team Totals, 2000-2005 NFL Seasons,




Figure 6: Comparison of NFL Player Cap Share of Team Total by Cap Rank to a Pareto Model and a In-In Model.

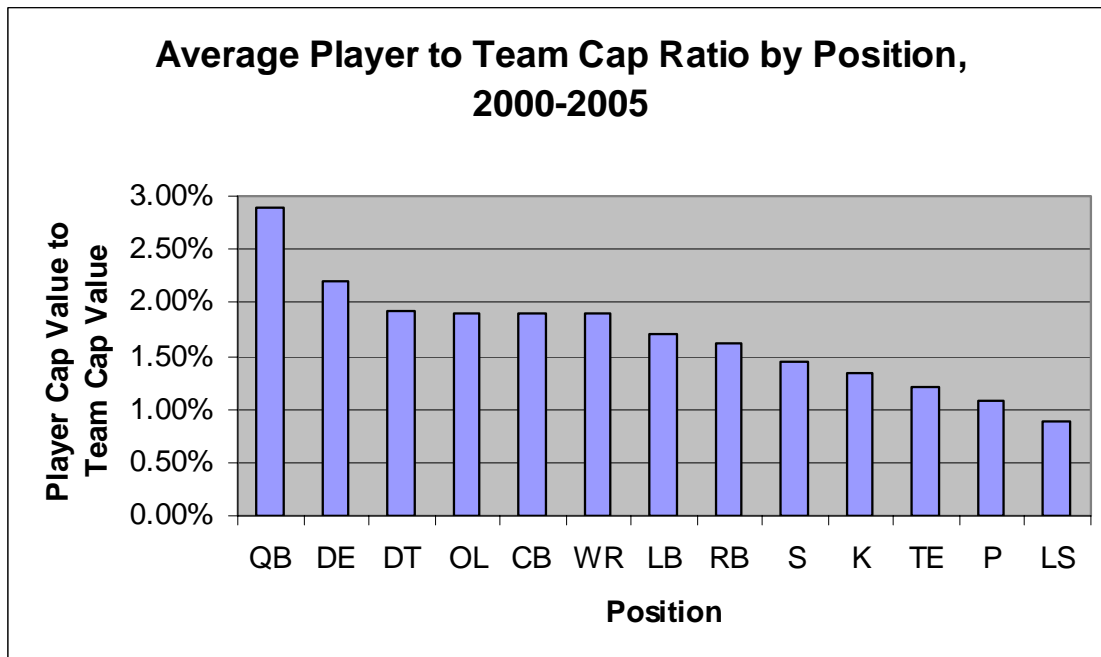


Figure 7: Mean Player Share of Team Cap Spending by Position, NFL, 2000-2005.




Figure 8: Mean Player Share of Team Cap Spending by Player Rank on Team and Playing Position (Special Teams Players Omitted).




Figure 9: Top 150 Career QB Passer Ratings for NFL Seasons Played During Salary Cap Era, Minimum 50 Pass Attempts.




Figure 10: Top 250 Career RB Rushing Averages for NFL Seasons Played During Salary Cap Era, Minimum 100 Rushing Attempts.

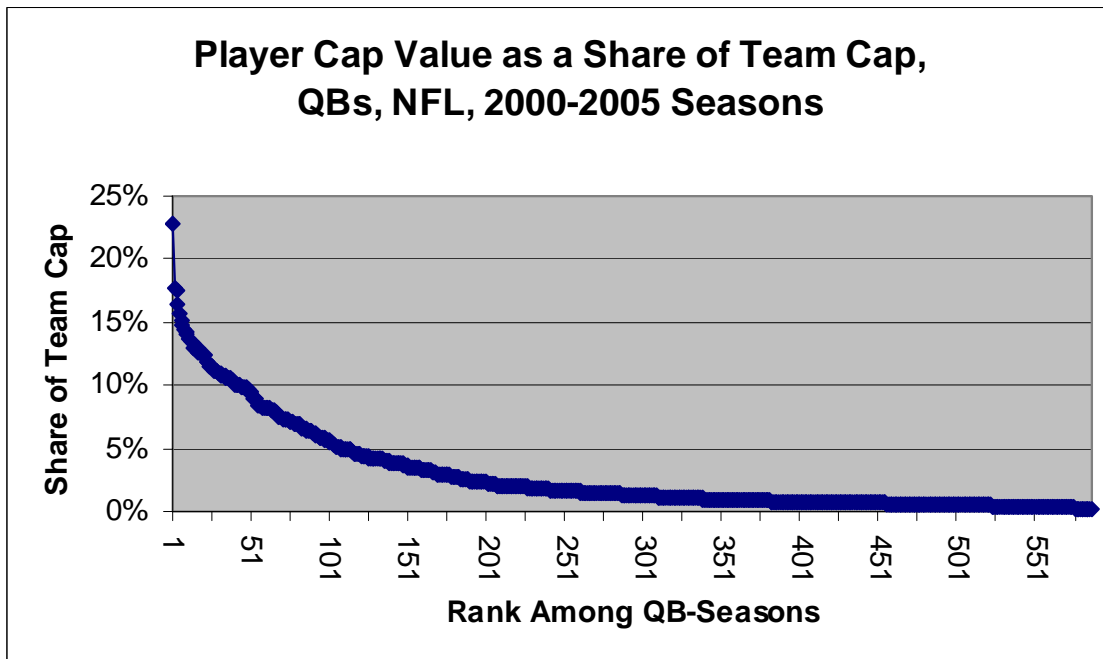


Figure 11: Player-Season Cap Value as Share of Team Cap by Rank of Player-Season Among All QB Player-Seasons, 2000-2005.




Figure 12: Player-Season Cap Value as Share of Team Cap by Rank of Player-Season Among All RB Player-Seasons, 2000-2005.

	2000	2001	2002	2003	2004	2005
2000	1.000	0.573	0.090	0.262	0.212	-0.181
2001		1.000	0.382	0.221	-0.019	-0.347
2002			1.000	0.226	-0.200	-0.502
2003				1.000	0.387	0.030
2004					1.000	0.131
2005						1.000

Table 2: Correlation Matrix for Team Cap Ratio to League Mean, 2000-2005 Seasons




Figure 13: Team Season Winning Percentage vs. Team Cap Spending, 2000-2005 NFL Seasons




Figure 14: Team Season Win Percentage vs. Gini Coefficients for Team Cap Spending Distribution for NFL Teams, 2000-2005 NFL Seasons




Figure 15: Mean Ratio of Winning Teams' to Losing Teams' Player Cap Spending by Player Rank, 2000-2005 NFL Seasons

Constant	0.240 (0.969)	0.119 (0.638)	0.0548 (0.408)
Lag(WinPct)	0.129 (1.51)	0.161 (2.08)**	0.162 (2.17)**
Ratio Cap Spending to League	0.305 (1.72)*	0.391 (2.66)**	0.366 (2.67)**
Lag(Cap Spending Ratio)	-0.051 (-0.296)		
Gini	0.315	-0.362 (-0.569)	
Lag(Gini)	-0.545 (-0.611)		
R²	0.060	0.075	0.076
Adj R²	0.027	0.059	0.066
F	1.79	4.72**	7.65**
N	146	180	189

Note: * indicates significance at the 10% level; ** indicates significance at the 5% level.

Table 3: OLS Regressions, Dependent Variable = Team Season Win Percentage

References

Cox, R.A.K. and J.R. Hill. 2005. The effect of the new NBA agreement on superstar pay. Accessed 5 June 2007 from <http://www.westga.edu/~bquest/2005/nba/NBA1.htm>.

Chung, K.H., and R.A.K. Cox. 1994. A stochastic model of superstardom: an application of the Yule distribution. *The Review of Economics and Statistics*, v.76, no.4, pp. 771-775.

Forbes. 2006. NFL team valuations. (31 August). Accessed 6 June 2007 from http://www.forbes.com/lists/2006/30/06nfl_NFL-Team-Valuations_Rank.html.

Hall, S., S. Szymanski, and A.S. Zimbalist. 2002. Testing causality between team performance and payroll: the cases of Major League Baseball and English soccer. *Journal of Sports Economics*, v.3 no.2, pp. 149-168.

Kahn, L.M., and P.D. Sherer. 1988. Racial differences in professional basketball players' compensation. *Journal of Labor Economics*, v.6, no.1 , pp. 40-61.

Kesenne, S. 1996. League management in professional sports with win maximizing clubs. *European Journal for Sport Management*, v.2, no.2, pp. 4-22.

Kowalewski, S. and M. Leeds. 1999. The impact of free agency and the salary cap on the distribution and structure of salaries in the National Football League. In J. Fizel, E.

Gustafson, and L. Hadley, *Sports Economics: Current Research*, pp. 214-225.
(Westport, CT: Praeger).

Larsen, A., A.J. Fenn, and E.L. Spenner. 2006. The impact of free agency and the salary cap on competitive balance in the National Football League. *Journal of Sports Economics*, v.7, no.4, pp. 374-390.

Lawrence, M. (n.d.) Football 101: the salary cap. Salary Cap Room and the Draft. Accessed 6 June 2007 from <http://football.calsci.com/SalaryCap3.html>.

Lazear, E. and S. Rosen. 1981. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, v. 89, no.5, pp. 841-864.

Leeds, M. and S. Kowalewski. 2001. Winner take all in the NFL: the effect of the salary cap and free agency on the compensation of skill position players. *Journal of Sports Economics*, v.2, no.3, pp. 244-256.

Pareto, V. 1897. *Cours d'Economie Politique*, Volume ii. (Lausanne: F. Rouge).

Porter, P.K., and G.W. Scully. 1996. The distribution of earnings and the rules of the game. *Southern Economic Journal*, v.63, no.1, pp. 149-162.

- Quirk, J. and R. Fort. 1992. *Pay Dirt: The Business of Professional Team Sports*. (Princeton, NJ: Princeton University Press).
- Romer, D. 2006. Do firms maximize? Evidence from professional football." *Journal of Political Economy* 114(3).
- Rosen. S. 1981. The economics of superstars. *American Economic Review* 71(5).
- Scully, G. 1974. Pay and performance in Major League Baseball. *American Economic Review*, v.64, no.6, pp. 915-930.
- Simon, H. 1955. On a class of skew distribution functions. *Biometrika* 42.
- Sloane, P. 1971. The economics of professional football: the football club as a utility maximizer. *Scottish Journal of Political Economy*, v.17, no.2 , pp. 121-146.
- Spofford, M. 2006. Packers rise to no. 7 in NFL revenue rankings. (17 June). Accessed 6 June 2007 from <http://www.packers.com/news/stories/2006/06/17/1/>.
- USA Today. 2006. Salaries database. Accessed 18 September 2007 from <http://asp.usatoday.com/sports/football/nfl/salaries/default.aspx?Loc=Vanity>.

Yule. 1924. A mathematical theory of evolution, based on the conclusions of Dr. J.C. Willis, F.R.S.," *Philosophical Transactions of the Royal Society B* B213.

Notes

¹The most recent NFL agreement specifies that this value should be 59.2%, and includes some additional revenue transfers from higher-local revenue teams to lower-local revenue teams.

² Teams may also carry up to 8 first- or second-year “developmental” players that can be elevated to the main roster or signed by other teams. This group is colloquially called the “taxi squad,” apparently because in the 1950s, legendary Cleveland Browns coach Paul Brown would find local cab driver jobs for a *de facto* reserve pool of players who did not make his roster.

³ Due to injuries and other factors, teams generally had several more players on its roster during each season than the limit of 53, however, a few teams had fewer than 55 players in some seasons.

⁴ The Pareto probability distribution shown in Figure 6 is $f(x, a, k) = (ak^a)/(x^{a+1})$, where a is estimated to be 0.00274 and k is estimated in Figure 5 to be 0.3000. However, a Cramer-von Mises (W4) empirical distribution test performed using E-views rejects the null hypothesis that the distribution is strictly Pareto at the 5% level. The constant, coefficient, and R^2 of the ln-ln model are -1.4878, -0.9642, and 0.949, respectively.

⁵ Also note that this model neglects the fact that in the NFL that on each roster there are a subset of players who are monopsonized via the transfer restrictions imposed by the collective bargaining agreement with the players. Each year, a portion of each team’s cap is designated as a rookie salary cap for drafted players; i.e., a cap within a cap; the value depends on the number and position of the team’s draft picks prior to that season. Nearly all teams use all of their rookie cap allotment, which averages about 4% of the total cap, spread across about the 10% of the players on each team who are rookies.

⁶ Note that a sufficient, but not necessary, condition for the existence of a solution to the optimization problem is that $(\partial^2 U / \partial W^2)(\partial^2 W / \partial q_i^2)(\partial^2 q_i / \partial c_i^2) < 0$.

⁷ An augmented Dickey-Fuller test on the data series shown in Figure 14 fails to reject a unit root, although the presence of a unit root is rejected for the first-differenced series. The first difference of this series exhibits an ARMA(2, 2) structure ($R = 0.379$). Further analysis of the significance of this characteristic is left for future research.