A Power Law Distribution for Tenure Lengths of Sports Managers

Toke S. Aidt[*1], Bernard Leong[#] , Daniel Sgroi[*] and William C. Saslaw[$]

[*]Faculty of Economics, University of Cambridge, U.K.

[#]Department of Bioinformatics, Wellcome Trust Sanger Institute. U.K.

[$]Department of Astronomy, University of Virginia, Charlottesville, VA, U.S.A.

Abstract

We show that the tenure lengths for managers of sport teams follow a power law distribution with an exponent between 2 and 3. We develop a simple theoretical model which replicates this result. The model demonstrates that the empirical phenomenon can be understood as the macroscopic outcome of pair-wise interactions among managers in a league, threshold effects in managerial performance evaluation, competitive market forces, and luck at the microscopic level.

Keywords: Power law distributions, complexity, tenure of managers, managerial turnover, competitive sports.

PACS codes:

---

[1] Corresponding author: Faculty of Economics, University of Cambridge, Cambridge CB3 9DD, United Kingdom.

E-mail: toke.aidt@econ.cam.ac.uk. Phone: +44 1223 335231. Fax: +44 1223 335475.

1. Introduction

The ideas of complex systems, originating in physics [1,2,3], are expanding into the social sciences. In economics the approach gives insights into patterns of regional specialization [4,5], economic fluctuations [6], trading activities in financial markets [7], Zipf's Law of cities [8] and other social phenomena. The science of complex systems maintains that macroscopic properties can emerge from large numbers of relatively simple microscopic interactions. This is in contrast to (neoclassical) economic theory, which builds on deductive analytical methods. An approach based on concepts arising in complex systems therefore provides new ways of viewing old questions [9].

This paper makes two related contributions towards understanding complex systems. First, we present new evidence that the distributions of tenure lengths for managers of sports teams, including football (soccer), American football, and baseball in many countries obey power laws. These have exponents between 2 and 3. Power laws have previously been observed in the sizes of firms and cities [8,10], gross domestic product per capita [11], turnover of rulers [12], wealth distribution [13] and stock market returns [14]. Stephen Durlauf [15] surveys these results in detail. Second, we propose a computable model that replicates the empirical power laws for the tenure of sports managers. The model demonstrates that the empirical phenomenon can be understood as the macroscopic outcome of pair-wise interactions among managers in a league, threshold effects in managerial performance evaluation, competitive market forces, and luck at the microscopic level. A surprising implication is that factors, such as talent, effort choices, and selection and matching processes that normally play a role in economic models of tenure, are not essential for understanding the dynamic evolution of hiring and firing in competitive sports.

2. The Evidence

Our data records information on the tenure length of managers of local football (soccer) clubs in England, Switzerland, France, Spain, and Germany, as well as of national teams, of baseball clubs in the U.S.A. (both the National and the American League) and Japan, and of American football clubs (both the NFC and the AFC). The data set contains 7183 managers covering up to 130 years.[2] Our goal is to characterize the empirical distribution functions of manager tenure lengths. We find that these tenure distributions are essentially power law probabilities $p(t) = b_i t^{a_i}$ where t is tenure length in years, a_i is a sports-specific parameter reported below and b_i is a normalizing constant. Power laws are, in the context of complex systems, emergent properties. Therefore identifying such distributions may indicate the empirical relevance of ideas developed for stochastic complex systems.

In Figure 1a-e, we show the distribution of tenure lengths for managers in English football (soccer) clubs plotted logarithmically such that power laws are straight lines. The diagrams show the evolution over time: Figure 1a covers the period before 1900, Figures1b-1d cover successive periods and Figure 1e covers the period from 1874 until the present. We clearly see that the distribution *evolves*, and a power law emerges. We find a similar emerging pattern by just dividing the sample into four equal sized periods (not shown). This rules out the possibility that the pattern emerges purely because we add more data points. However, the number of managers fired within the first year of employment falls below the power law: too few managers are fired. Firing decisions within the first year (season) are likely to be determined by very different processes than decisions later in a manager’s career. When first hired, managers generally have a “honeymoon” period to build a team, so they are not really judged until their second season.

[2] The data material is available upon request from the corresponding author.

Moreover, a typical contract includes clauses that make it expensive to fire a manager within a season; most market activity occurs between seasons. Our data is coded by allocating managers fired between the first and second year to the one year bin, managers fired between the second and third year to the two year bin, and so on. Since managers are rarely fired on the day they are hired, this reduces the number of recorded firings allocated to the zero year bin. For these reasons, we focus on the subset of managers with tenures greater than one year. We also ignore managers with exceptionally long tenures (more than 25 years).

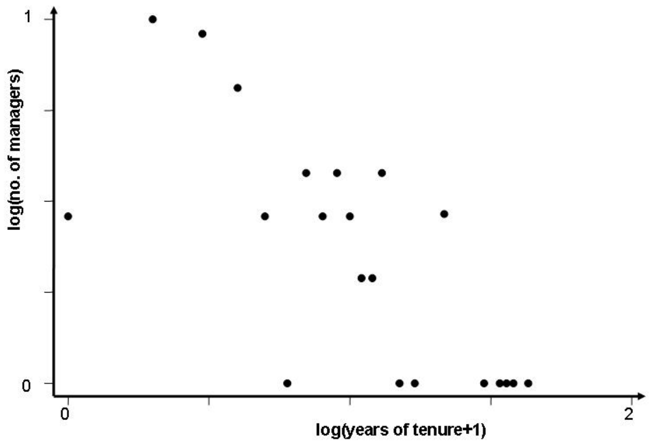

Figure 1a: The distribution of tenure length for English football managers, 1874-1900

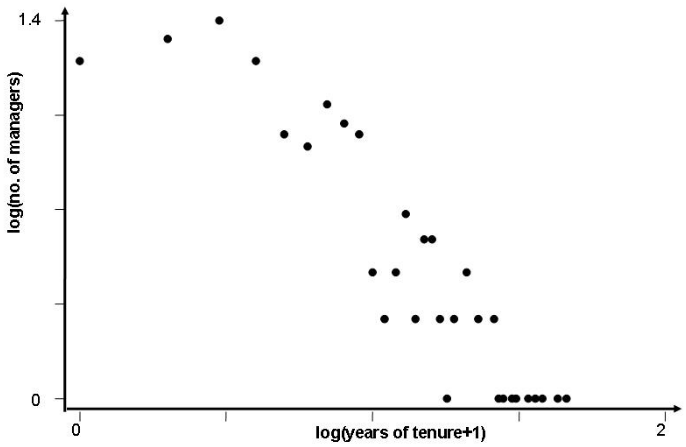

Figure 1b: The distribution of tenure length for English football managers, 1874-1920

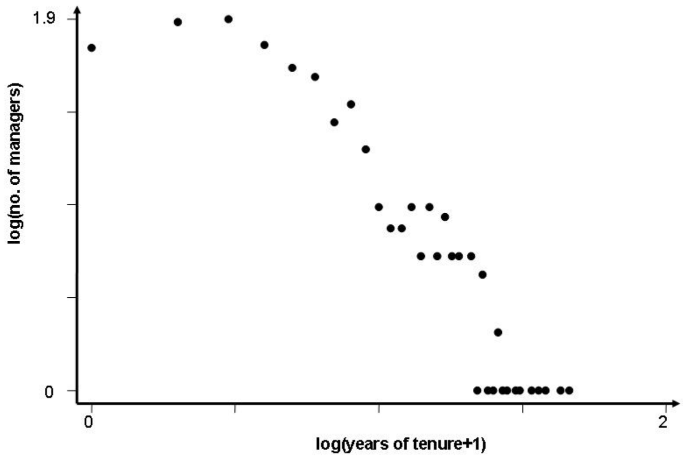

Figure 1c: The distribution of tenure length for English football managers, 1874-1960.

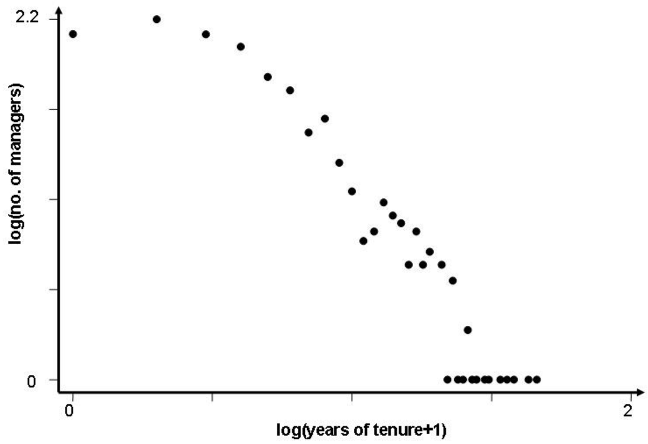

Figure 1d: The distribution of tenure length for English football managers, 1874-1980.

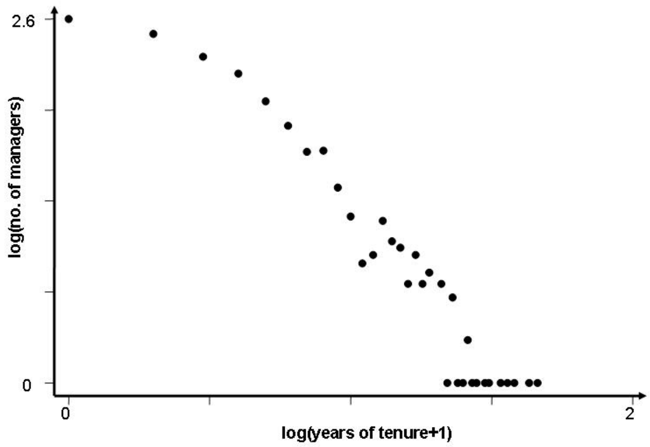

Figure 1e: The distribution of tenure length for English football managers, 1874-2005.

Figure 1: The distribution of tenure lengths for English football mangers, successive periods, 1874 to 2005.

Table 1 reports the results of a formal test of the hypothesis that the tenure length of managers in these sports clubs can be described by a power law probability distribution and, if so, reports an estimate of the parameter a_i of the empirical distribution function. We test for a linear relationship between the logarithm of the frequency of tenure length and the logarithm of tenure length itself. This is done by fitting an Ordinary Least Squares regression line to the frequency data and using a Ramsey regression specification error test (RESET) for omitted variables to test for linearity. The RESET tests if powers of the logarithm of tenure length can improve the fit of a simple linear regression. Rejection of the null hypothesis that these powers belong in the regression is statistical evidence for a linear relationship.

Inspection of the p-values from the RESET, reported in the last column of the table, shows that the data supports (at the 95 per cent level of confidence) that the tenure lengths of sports managers follow power law probability distributions in all cases except for football (soccer) managers in France. The estimated exponent of the power law reported in column 4 ranges from -2.23 for football managers in England to -2.98 in Switzerland and is surprisingly stable across sports. These results are also robust to organizing the data in two year bins.

3. A Theoretical Model of Manager Turnover in Competitive Sports

Managers of sports teams work in a very competitive environment with a clear-cut criterion for success or failure: the performance of a manager is judged by the performance of the team he manages. Roughly speaking, for each match his "reputation" is either enhanced or diminished depending on the result, and it is natural to suppose that a manager is replaced when his accumulated "reputation" reaches a lower threshold, or he is poached by another club when his accumulated "reputation" reaches an upper threshold. We propose a computational model that, under a range of circumstances, can replicate the empirical tenure length distributions. We model tenure decisions in a sports league as the outcome of the pairwise interaction of managers in a league, competitive market forces, threshold effects in the evaluation of managerial performance, and luck. The model is deliberately designed to be simple and it omits many factors that could affect tenure length. This include skill/talent differentials between managers, the importance of particular matches, the matching process by which managers and teams are paired, and the impact of the sequence of results on firing decisions. Our objective is to isolate the minimum number of factors needed to understand the observed empirical regularity.

The core of the model is a simple round-robin tournament similar to the premier league in England.[3] This links the managers in a simple network. The round-robin tournament works as follows: (i) there are twenty managers and teams in a league; (ii) each manager (team) plays two games against the other managers (teams), so that there are 38 games in total, with 19 "home" and 19 "away" games in each tournament; and (iii) in each game, the probabilities of win, lose and draw are 37:26:37. These probabilities correspond to those observed in the English football league over the period 1881-1991 [16] and are assumed to be independent of the identity of the managers involved in the match.

The turnover of managers is governed by the following rules that relate tenure length to performance in the tournament. First, the model is initialized with twenty randomly selected managers, each with two attributes, *reputation* and *tenure*. Each manager's reputation when the season starts is a random positive integer with a constant probability distribution between 0 and the poaching threshold (see below) and each manager starts with a random tenure length with a flat distribution between 1 and 40 years. Each manager gains reputation every time his team wins a game and loses reputation when his team loses. In our baseline specification, we employ a symmetric scoring rule, which assigns +2 to the manager's reputation if he wins a game, -2 if he loses and 0 if he draws. The scoring rule encapsulates how the board of directors, fans of the club, the media and others judge the performance of the manager. Variations in the scoring rule correspond to changes in the mapping from performance to valuation. We assume that the values assigned to a win, lose or draw by the scoring rule are the same for all managers and matches. A manager's tenure depends on the evolution of his reputation, and he stays in office while his reputation remains within lower and upper bounds; otherwise he is replaced. Termination of tenure can happen for four reasons: (i) the manager is fired for poor performance when his

[3] The source code for the model is available upon request from the corresponding author.

reputation falls below the lower threshold, called the *firing threshold*; (ii) the manager is poached by another club as a result of good performance when his reputation reaches the upper threshold, called the *poaching threshold*; (iii) the manager retires due to old age if his tenure exceeds a fixed number of years, called the *ossification threshold*; (iv) the manager's team is relegated (to another lower league) because it has the lowest reputation in the league at the end of the season. When a manager is either fired, his team relegated or he retires his place in the league is taken by another manager with tenure length of zero and a reputation selected at random from a flat distribution between 0 and the poaching threshold. We set the firing threshold at zero and maximum tenure length at 25 years. The important free parameter of the model is the poaching threshold. For each set of parameter values, we simulate the tournament repeatedly and for each run, we record the distribution of tenure lengths corresponding to a hundred years.

4. Replicating the Empirical Distribution Functions

This section addresses two questions. Firstly, can the model produce a power law probability distribution of tenure lengths? Secondly, will an appropriate choice of the poaching threshold replicate the observed data from the five national football leagues?

To answer these questions, we examine the normalized frequency distributions of tenure lengths generated by the model and compare them with the observed data. The answer to the first question is affirmative: for a very broad range of parameter values within the class of predetermined, fixed and symmetric scoring rules, the model produces a tenure length distribution that is statistically indistinguishable from a power law distribution. Figure 2 reports one example of the normalized tenure length distribution with a poaching threshold of 35. Similar results obtain for other probability distributions of win, draw, or lose. The assumption of a symmetric scoring rule is important for this result. Asymmetric rules where, say, a win increases the reputation much more than a loss reduces it do not lead to power law distributions (not shown). It is also important to use the same symmetric scoring rule for every match. When the values allocated for win, draw or lose are different for each match, the model fails to generate power law distributions of tenure lengths. An important requirement for a power law distribution is therefore that all matches are equally important, both over time and in symmetrizing their reward/penalty ratios.

The answer to the second question is also affirmative and similar results can be obtained for the other sports. Figure 2 reports the actual (normalized) tenure length distribution for managers in the five national football (soccer) leagues. We observe that with the chosen parameter values, the model replicates this distribution almost perfectly. These parameter values, however, need to be chosen carefully to replicate the data. For example, decreasing the poaching threshold steepens the distribution function, as more frequent poaching leads to more frequent manager turnover. Likewise, smaller rewards and penalties for wining and losing, say +1 and –1 instead of +2 and –2, flattens the distribution function, since then managers usually stay in office for longer.

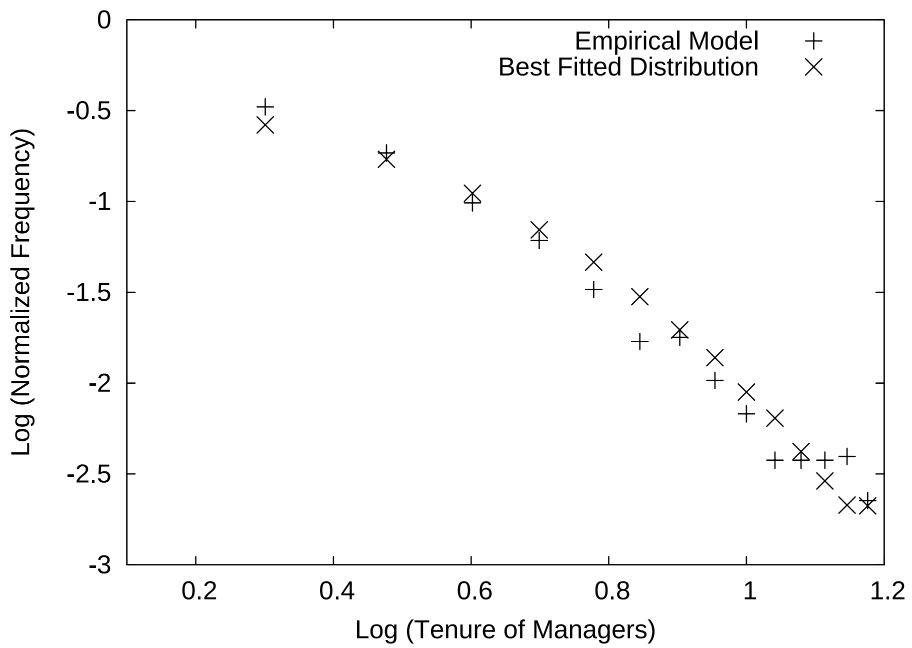

Figure 2: Comparison of the observed tenure distribution for all football (soccer) managers and the best fitting model-generated distribution function.

5. Discussion

Per Bak [2] and others have postulated that power laws are a statistical consequence of complex systems that are capable of self-organization at critical points. To what extent can the empirical regularity reported above be related to self-organized criticality? A priori there are reasons why a social organization such as a sports league might evolve to a critical state. Social organizations gradually evolve in response to internal and external economic, legal and social pressures. Decisions within social organizations result from complex interactions among many factors. Occasionally, a major event such as a managerial or other reorganization can be triggered. However, for a social organization to evolve into a self-organized critical state, the organization must have the ability *endogenously* to select the particular parameter values that make it critical and have a mechanism that prolongs this state. Usually such mechanisms are associated with three requirements: an input (of e.g., particles, energy, or in our case of managers), a mechanism for maintaining a meta-stable state (e.g., a finite and fixed range of threshold values for replacing managers), and a form of dissipation (e.g., the actual replacement of managers through firing, poaching or retiring). However, these mechanisms do not guarantee self-organized criticality. Our model requires given *exogenously* fixed parameter values to catalyze re-organization which can be characterized by a power law distribution.

Yet, going beyond the model there are a number of additional factors that point in the direction of self-organized criticality. First, competitive market forces might bring sports leagues to a point of maximum efficiency related to self-organized criticality. A team wishes to maximise its profit. This means maintaining attendance at games, together with maximising television and advertising revenues. A key feature is making each game important, as fans and advertisers alike will not wish to financially support irrelevant games. We could imagine that rewards and penalties might initially be distributed via a probability distribution with high variance, which allows the rewards from some games to be so low as to make their outcome irrelevant. Hence, we would expect self-organisation to evolve the system to one with a low variance of rewards, pushing the system to a point distribution. This is further supported by the fact that our model can match the empirical power law distributions when all matches have the same reward structure, but is unable to do so when the reward structure for each match is randomly determined.

Second, competition in a sports league is a zero sum game where what matters is relative, not absolute performance. Managers, therefore, finds themselves involved in an arms race and subject to the Red Queen principle that the more things change, the more they stay the same [17]. The Red Queen principle, although not sufficient to generate the dynamics associated with a system in a self-organized critical state [9], is indicative of why sports leagues once they have reached a self-organized critical state might stay there for a prolonged period of time.

Although our evidence is specific to team sports, we believe that similar tenure distributions would obtain for other social organizations such as large corporations. This conjecture is further supported by the fact that the tenure lengths of political leaders (Popes, Roman Emperors, British Prime Ministers etc.) follow power law distributions, albeit with lower exponents than those found for sports managers [12]. The power law exponents for various classes of political leaders are influenced by their method of selection as well as of termination.

Our results suggest, perhaps surprisingly, that innate ability need not play a major role in success.

1. P. Bak, C. Tang, and K. Wiessenfeld, *Phys. Rev. A* 38, 364-74 (1988).
2. P. Bak, *How Nature Works: The Science of Self-Organised Criticality* (Copernicus, New York, 1996).

3. H.J. Jensen, *Self-organized Criticality. Emergent Complex Behaviour in Physical and Biological Systems* (Cambridge University Press, Cambridge, U.K., 1998).

4. T. Schelling, *Micromotives and Macrobehaviour* (W.W. Norton, New York, 1978).
5. P. Krugman, *The Self-organising Economy* (Blackwell Publishers, Oxford, U.K., 1996).
6. H.E. Stanley, P. Gopikrishnan, V. Plerou, *Physica* A 283(1-2), 31-41 (2000).
7. X. Gabaix, P. Gopikrishnan, V. Plerou, H.E. Stanley, *Nature* 423, 267-270 (2003).
8. X. Gabaix, *Quarterly Journal of Economic* 114(3), 738-67 (1999).
9. S.M. Markose, *Economic Journal* 115, F159-F192 (2005).
10. R. Axtell, *Science* 293, 1818-20 (2001).
11. D. Canning, L. Amaral, Y. Lee, M. Meyer, H.E. Stanley, *Economics Letters* 60, 335-41 (1998).
12. W.C Saslaw, R Tuck, mimeo (2002).
13. V. Pareto, *Cours d'Economique Politique* (Macmillan, London, 1897).
14. T. Lux, M. Marchesi, *Nature* 397, 498-500 (1999).
15. S.N. Durlauf, Complexity and empirical economics. *Economic Journal* 115, F225-F243 (2005).
16. Rothman's Football Yearbook, (London, 1970, 1980, 1991, 2000).
17. P. Anderson, Local moments and localized states. Nobel Lecture http://www.nobel.se/physics/laureates/1977/anderson-lecture.pdf (1977).

Supplementary Information includes the data set and a description of its sources. The source code for the computer simulation is available upon request.

Acknowledgements

We thank James Norris and Thomas Down for helpful suggestions, Laura Webster for excellent research assistance, and the Faculty of Economics, University of Cambridge for financial support. An International Fellowship from the Agency of Science, Technology and Research, Singapore supported Bernard Leong.

Table 1: The Estimated Parameter of the Tenure Length Distribution and the Results of the RESET in Different Sports and Countries

Sport	Country	Period	a_i	RESET[1]
Football	England	1874-2005	-2.23	0.13
Football	Switzerland	1921-2002	-2.93	0.12
Football	France	1928-2003	n.a	0.03
Football	Spain	1910-2003	-2.75	0.29
Football	Germany	1947-2003	-2.51	0.08
Football	All league managers	1874-2005	-2.75	0.23
Football	National teams	1902-2003	-2.92	0.14
Baseball	USA	1875-2003	-2.63	0.14
Baseball	Japan	1936-2003	-2.41	0.65
Am. football	USA	1920-2003	-2.08	0.14

Note: [1] The p-value of the RESET (H_0: The model is not linear).