TI 2017-116/VII
Tinbergen Institute Discussion Paper

The inefficient advantage of experience in the market for football managers

Thomas (T.L.P.R.) Peeters[1]
Stefan Szymanski[2]
Marko Terviö[3]

1: Erasmus School of Economics, ERIM; Tinbergen Institute, The Netherlands
2: University of Michigan
3: Aalto University

The inefficient advantage of experience in the market for football managers

Thomas Peeters[1]

Erasmus School of Economics, Tinbergen Institute and ERIM

Stefan Szymanski

University of Michigan

Marko Terviö

Aalto University

Abstract

We study hiring in a labor market where worker ability can only be observed on-the-job, but quickly becomes public information after labor market entry. We show that firms in these markets have a socially inefficient incentive to hire low talented, experienced workers instead of more promising labor market entrants, either when an extremely poor hire may bankrupt the firm, or when workers cannot commit to long-term contracts. In a dataset covering 38 years of hiring in the English labor market for football managers, we find that in around one quarter of all cases, where a firm hires an experienced worker, this experienced worker has an estimated ability below the average ability of recent labor market entrants. We argue this hiring behavior is inefficient, because it has persistently depressed the average ability of the active manager labor force over our sample period.

JEL Codes: M51, J63, J24, Z22

Keywords: hiring, labor market entrants, worker ability, European football

[1] Corresponding author. E-mail: peeters@ese.eur.nl. Phone: +32 494 12 49 36. Postal: Erasmus School of Economics, Postbus 1738, 3000 DR Rotterdam, Netherlands. We benefited from comments by Jan van Ours, Ignacio Palacios-Huerta and Arne Warnke. We further thank seminar and conference participants at the University of Sankt Gallen, Imperial College Business School, DIW Berlin, Erasmus School of Economics, Universitad Autonoma de Barcelona, Prague School of Economics, University of Reading, 2016 Colloquium on Personnel Economics in Aachen, 2016 European Conference on Sports Economics in Groningen and 2017 European Association for Labour Economics meeting in Sankt Gallen.

1 Introduction

Since the work of Jovanovic (1979) it has been well known that the option value of employment relationships is critical for efficient turnover decisions in labor markets. Hiring a worker who is known to be of average ability, is socially inefficient if it is possible to instead hire someone who is equally able in expectation, but could turn out better or worse. The basic intuition is that successful workers will be retained for longer than unsuccessful ones, so upside risk is more valuable than downside risk, and hence riskier candidates should be selected over safer candidates. Terviö (2009) argued that when worker performance is publicly observable, as in creative industries such as music and sports, the finders do not get to keep the stars they discover---at least not at a wage that would let them earn rents to compensate for the initial investment. When differences in ability entail large differences in economic value, it is unrealistic to expect entry-level workers to pay firms for the full value of the chance to be discovered. It is then rational for employers to take a short-term view in their hiring decisions, which implies a lenient standard for hiring experienced workers and results in inefficiently low average ability in the industry.

In this paper, we argue that this inefficiency can be expected and indeed seen in the labor market of professional football managers. We first show how liquidity constraints on the part of firms may provide an additional explanation for this phenomenon. We then use comprehensive data on managerial careers in the English football industry over several decades to estimate the expected ability of experienced managers and entrants. We find that it is not uncommon for an experienced manager with below-average revealed ability to be hired over entrants with higher expected ability. This is actually worse than just ignoring the upside potential of entrant managers, and cannot be explained within the earlier models, but is consistent with our model if firms face a risk of insolvency.

The information structure in the market for football managers is not well adapted to socially efficient hiring. For managers on-the-job, there is a clearly observable and common benchmark of success, winning games, and publicly available financial statements provide reliable information on the resources available, which vary considerably among clubs. In our analysis, we define managerial ability as a manager’s measured contribution to team winning, after controlling for inputs. Our econometric method enables us to recover individual fixed effects representing the manager’s ability, which we identify with the public assessment of that ability. We adopt a rolling window estimation procedure to capture changes in perceived managerial ability in reaction to observed success and failure. Our results indicate that observable managerial characteristics at the time of market entry are of little use in predicting ability, but repeated observations of performance quickly enable precise predictions of a manager’s future ability estimates as his career progresses. This suggests that it is very hard to predict the ability of a football manager before he gets a chance to manage a professional football club, but once employed, performance is quickly revealed to everyone, including competing employers.

The key market imperfection in Terviö (2009) is that entering workers are liquidity constrained, otherwise they could pay employers for the opportunity to find out their ability. However, in English football it is apparent that the employers themselves are also liquidity constrained. Among the approximately 100 football clubs in our data, there have been 67 instances of legal insolvency proceedings in the last quarter century (see Beech et al (2010) and Szymanski (2017)). Szymanski (2017) showed that insolvency is usually associated with underperformance in the league competition, which in turn can partly be attributed to the ability of the manager. We therefore extend the model of Terviö (2009) to allow for liquidity constrained firms. In the model, a sufficiently low ability manager can drive a firm into insolvency, which means that its current owners would lose their stake in the club. It turns out that the owners can then be better off by hiring an experienced manager with a low revealed ability, rather than risking an inexperienced entrant whose ability is higher in expectation but could turn out to be even lower. Our model thus predicts that firms may choose to hire experienced managers, dubbed "substandard" managers, who are even worse than the average inexperienced manager.

Our dataset contains information on approximately 75,000 games played in the four divisions of English professional football over the period 1974-2011. A typical league season comprises about forty games for each team. We retrieve financial information, such as revenue, wages, profitability from audited accounts filed with Companies House. We combine this with a database of personal characteristics of football managers (including nationality, age, experience and previous career as a player) for the population of around 950 managers who were active in this period.

We empirically study the efficiency of hiring decisions by comparing the ability of rehired managers with a counterfactual of hiring a market entrant. To control for input use in the estimation of managerial abilities, we use individual effects estimators à la Abowd, Kramarz, and Margolis (1999), whose method was developed in order to identify the contribution of specific individuals in the context of team production. In the same way, while football represents a classic example of team production, we are able to identify the contribution of a specific individual, namely the manager. We estimate the model using only observations up to the time of hiring, hence our ability measure is consistent with what was known by potential employers at each point during a manager's career. Our measure of ability for incumbent managers is straightforward: we use our best estimate of their expected ability based on data that was publicly available at the time of their rehiring. The distribution of ability for the potential entrants that were not hired in their stead cannot be directly observed and so we use the estimated ability of actual entrants in the previous five years.

We find that about one quarter of all rehired experienced managers are of lower ability than the average entrant at the time of hiring. Clubs in lower divisions, which are less productive in generating revenues out of wins, make more substandard rehires than their more productive rivals. Nevertheless, employers in all divisions use the information generated by games played. In comparison to their more able counterparts, low ability managers are more likely to be fired and, upon losing their job, less likely to be rehired by another club and less likely to start their next job in a higher division, where more money is at stake. These effects are stronger for low ability managers who have coached more games, and whose ability can be inferred more precisely. Moreover, learning about manager ability is fast – within two months past performance can account for upward of 80% of current performance. A simple counterfactual exercise suggests that the hiring policies we observe in the data have depressed the worker ability of the active labor force of football managers in England.

Previous studies of inefficient hiring have used data from online labor markets. A field experiment conducted by Pallais (2014) showed that there is inefficiently little hiring of inexperienced data entry specialists: employers provide public reviews of worker performance, but do not take into account the value of this information in later hiring decisions. Stanton and Thomas (2014) showed that intermediaries provide a partial solution to this problem by producing quality signals about inexperienced workers. One benefit of our data is that we are able to track entire careers, some of which last several decades. Our data also concerns a job, which is both very high profile and deemed to be of great significance to consumers (fans) who buy the product. However, since we do not have a source of exogenous variation in hiring policies or in publicly available information, our estimates of inefficiency are necessarily more indirect.

The modern empirical literature on employer learning about worker ability began with Farber and Gibbons (1996), who studied wage determination in a setup where all employers observe worker performance and thus obtain the same increasingly accurate estimate of worker ability as the career progresses. The focus of the literature has been on explaining career wage dynamics, promotions, and inter-firm mobility, with extensions such as time-varying ability (Kahn and Lange, 2014) and firm heterogeneity (Abowd, Kramarz, and Roux, 2006). Our focus is quite different: we look at the selection of workers to a specialized occupation at the level of an entire industry. We do not have individual wage data, and the turnover we are interested in is in and out of the industry.

One central theme in the literature on employer learning is the role of asymmetric information. Schönberg (2007) finds a significant role for asymmetric information about worker ability between firms but only for college graduates, whereas Kahn (2013) finds it pervasive in the US labor market. While there is large uncertainty ex-ante about ability in the labor market for football managers, this uncertainty is likely to be symmetric. Without experience, neither the firm nor the would-be manager can know how well they are likely to perform. Ex post, the public nature of performance means the ability estimate of an experienced manager is public information. Thus, informational asymmetries between employers are unlikely to be an important factor. We assume that information is symmetric, which simplifies the analysis considerably.

Finally, there exists a substantial literature devoted to analyzing the labor market for football managers. One strand of this literature establishes that managers indeed exert a significant impact on the club's sporting performance, e.g. Frick and Simmons (2008), Frick et al (2010), Bridgewater et al (2011), Bell et al (2013) and Muehlheusser et al (2016). A second central topic in this literature entails the firing of football managers. There is a broad consensus that disappointing team performance is the main reason that managers get fired, see e.g. Audas et al (1999), Bachan et al (2008), Barros et al (2009), d'Addona and Kind (2014) and Van Ours and Van Tuijl (2016). The effect of replacing a manager on consequent performance is more ambiguous, with some authors documenting a (partial) positive effect (Tena and Forrest, 2007, and Schneemann et al, 2016), while others find no effect (Bruinshoofd and Ter Weel, 2003, Koning,2003, De Paolo and Scoppa, 2011, and Van Ours and Van Tuijl, 2016). As such, this literature has presented some evidence of excess firing at club-level. Boards may be too eager to fire managers, possibly because they over-attribute spells of bad results to low ability instead of bad luck or because the manager functions as a scapegoat. While we empirically confirm that clubs are more likely to fire a low ability manager, our analysis focuses primarily on hiring and firing at the level of the entire industry. Somewhat paradoxically, we find that the industry as a whole is too eager to re-employ experienced managers, even when individual clubs seem trigger-happy when it comes to firing the current manager.

In the remainder of this paper, we first sketch the institutional environment of the labor market for football managers in England. Next, we develop a theoretical model of hiring with firm liquidity constraints. Then we explain the dataset and empirical methodology we employ, followed by our main empirical results. In section 6, we present supporting empirical results, which examine the broader functioning of the labor market in our study. We review the main contributions of our work in the final section.

2 The market for football managers

A football manager is responsible for overseeing the training of players, managing tactics and directing the strategy of the team on and off the field. Managers further exert significant control over the hiring of players. All professional English football clubs are organized as limited liability companies with a board of directors and CEO responsible for the day-to-day management of the business, and it is very rare for a football manager to become a member of the board of directors. Hence, in purely business terms these individuals would typically be labeled middle managers. However, the public prominence of football means that many of them become household names. The activity of a football club is relatively simple and while commercialization has increased the financial size of the industry dramatically over the last quarter century, the functions of the manager have not changed significantly. The rules of the game itself have changed little, and winning games has always been the yardstick of managerial performance.

Most managers are drawn from the ranks of former players, and a few even become managers while still playing. Frequently a retired player joins the staff of a club and becomes engaged in activities such as training and scouting for playing talent. A career profile might involve running a junior team and acting as assistant to the manager before reaching the top job. However, high profile former players are often fast tracked through the system. Formal qualifications were not required before 2003, since when managers in the top division are expected to obtain a UEFA Pro License, which takes about a year of part-time study to acquire.[2] Exemptions from this requirement are however possible.

At club level, the median tenure of a football manager is about one year, but a significant portion of employment spells is much shorter. The 25th percentile of spell duration is just over 2 months. The board of directors can fire a manager whom they perceive to be unsuccessful. Dismissed managers may seek to re-enter the industry with another club, take employment at a lower organizational level (e.g. talent scouting) or move to a related business such as TV punditry. While average tenure is short, a managerial career usually spans several employment spells. As such, the median career length is slightly over two years, and only 23% of all managers have a career shorter than 1 year. At the other end of the spectrum, some managers have extremely long careers. The most famous example of recent years is (Sir) Alex Ferguson, who was manager of Manchester United from 1986 until 2013 and won 13 league championships, more than twice as many as the second best in the history of English football. Although exceptional, these long-lasting managers are naturally overrepresented at any one time in

[2] see https://www.thepfa.com/coaching/courses/qualifications

our game level data. For example, the 37% of managers with a career length under 2 years represent only 7% of all game observations.

The one hundred or so English professional football clubs are segregated into four hierarchical divisions of roughly equal numbers (currently 20 in the top division, and 24 in each of the 3 lower divisions). Each club plays every other club in its division (home and away) over a ten-month season. At the end of a season the worst performing N teams in a division are automatically sent down (relegated) to play in the division ranked below in the following season, to be replaced by the best performing N teams in that lower division (N has varied between 2 and 4 across divisions and time).

The consequences of relegation are severe. Relegation typically means lower attendance since your opponents are less attractive, and reduced access to broadcast revenues. In 2010, the average revenues in the four divisions were in proportion (from top to bottom) 25:4:2:1. Unlike the Major Leagues in America, there are limited mechanisms to promote revenue sharing and wage controls. Were rich clubs to agree to such mechanisms with their poorer rivals, they would only increase their own probability of relegation. Looked at from the other direction, the financial inequality between the divisions is a large incentive to devote resources to getting promoted to a higher tier.

While relative proportions have changed, over the last 40 years the structure of revenues and costs in football has not. Revenues are generated from four principal sources: match day (tickets sold, food and beverage, etc.), broadcasting, merchandising and sponsorship. The biggest change of recent years has been the extraordinary increase in the size of broadcast revenues, which for the top division went from less than £10 million per season in the mid-1980s to over £1000 million in 2011. Because broadcast revenues are concentrated in the top division this trend has increased inter-divisional inequalities. Producing commercial football games requires two essential assets: a stadium (which in England is mostly owned by, and paid for, by the club itself) and players. The market for players is highly competitive. There are large numbers of buyers and sellers, and, like managers, information on the ability of experienced players is widely available.

Larger clubs tend to command higher levels of support for any given level of success and hence have access to larger budgets. There is a substantial literature showing that teams with greater financial resources are more successful on the field, mainly because they can fund the acquisition of better players (see for example Szymanski and Smith, 1997 and Peeters and Szymanski ,2014). Large market teams therefore tend to move up to the higher divisions, which again increases the correlation between a club's divisional position and revenue potential. Football managers are naturally constrained by the resources under their command. Due to the divisional structure however, the variation in resources across teams that actually play each other is much smaller than the overall variation. In a given game, the ratio of the largest to smallest team payroll is on average 1.6 to 1, while the ratio of the average payroll in division one versus division four has increased from about 4 to 1 in the seventies to over 20 to 1 in recent years.

Insolvency is a relatively frequent event in English football, with an average of two to three legal insolvency events among the hundred clubs every year for the past quarter century (Beech at al, 2010 and, Szymanski, 2017). However, the clubs themselves almost always survive, usually being acquired by a local benefactor or a fund raising initiative organized by fans. Thus while there is turnover in the limited liability corporations that own the clubs, and turnover in the ownership of those companies, the identity of the football clubs is extremely stable. Insolvencies are more common in the lower divisions, because revenues are typically smaller and clubs may overreach themselves in pursuing promotion or experience significant revenue losses following relegation. The competitive pressure of the promotion and relegation system is also cited as a main contributor to the rapid turnover and low tenure among football managers.

3 A model of hiring with firm liquidity constraints

The previous literature on hiring with on-the-job learning has assumed that firms face no liquidity constraints. However, in the English football industry the threat of insolvency is an important concern (Szymanski, 2015 & 2017). In this section, we present a formal model that shows why we can expect liquidity constrained football clubs to be too eager to hire experienced managers, in a way that lowers the ability of managers active in the industry.

We use the model to examine how profit-maximizing football clubs hire managers, when they face a choice between inexperienced and experienced managers whose expected abilities may differ. There is always a pool of available inexperienced managers, whose ability will be a random draw from a known distribution, but the availability of experienced managers depends on past hiring decisions (and luck). We take it as a key feature of the industry that the ability of managers can only be inferred after observing job performance at a professional football club. The number of real job opportunities is scarce, and there will always be more potential good managers than could ever be employed as managers.

To understand the key economic trade-off here, notice that a managerial employment spell produces not just current output, but also information about a manager's ability, which is useful for future output. It is therefore socially inefficient to hire an experienced manager who is known to be only slightly better than a manager of unknown but expected average ability, because the latter may yet turn out to be substantially better than the average (Jovanovic, 1979 and MacDonald, 1988).

When worker performance is publicly observable, as it is for football managers, there are two theoretical reasons to suspect that the hiring policy that is optimal from the point of view of a profit-maximizing firm is not socially optimal. Both reasons lead us to expect a bias towards hiring experienced but so-so managers at the expense of new entrants.

The first problem that makes firms too lenient towards experienced managers is the fear of "poaching", which arises from the inability of managers to commit to long-term wage contracts. If successful managers are easily poached or have their wage quickly bid up by competing clubs, then a club that discovers a new star manager gains little from its gamble. This problem was analyzed in Terviö (2009) where it was shown that, in competitive equilibrium, a significant portion of jobs are populated by known mediocrities. If more novices were employed at the expense of these incumbents, this would result in higher average ability and higher welfare. The market imperfection behind this inefficiency is that novice managers (a) cannot commit to long-term wage contracts and (b) are liquidity constrained---otherwise they could "buy the firm" and then sell it to another novice if efficient.

The second problem arises when, not just workers, but firms face a liquidity constraint. This means that a bad "draw" from the manager ability distribution can cause the firm to make a loss that it is not able to cover, even if it could expect positive profits in the future. From the aggregate welfare point of view, a temporary setback of a very poor manager would be an acceptable loss when considered against the possibility of finding a very good one---but if the firm is not able to finance temporary losses, the future after a sufficiently bad manager has no value at all for the current owners. A liquidity constrained firm will therefore give more weight to the downside risk in managerial ability, and may hire a manager of known low quality over a more risky manager, even if the latter is expected to perform better. This is the key intuition in our model.

3.1 Model setup

Consider a firm with revenue equal to the ability of the manager. There is a population of untried managers with ability drawn from a known continuous distribution. Manager ability θ becomes known after one period of work, and managerial careers can last up to two periods. The firm is (potentially) infinitely lived and its objective is to maximize the expected present value of profits. It faces a simple liquidity constraint: revenue must exceed cost $c \geq 0$ in any period or else the firm goes bankrupt. The owners have limited liability, so period profits are $\pi(\theta) = \max\{\theta - c, 0\}$.

Here we assume that managers can commit to a two-period wage contract, so there is no poaching. The purpose is to analyze the effect of the firm liquidity constraint independently of commitment problems, which are already known to distort hiring. The number of clubs is scarce relative to the number of potential managers so novice managers are held to their outside wage, which we normalize at zero.

The firm then faces a real decision only in periods when it has an incumbent manager: should it hire a novice or retain the incumbent? This decision boils down to a rehiring threshold that we denote by ψ; the firm will retain managers when their ability is above ψ and otherwise hire a novice.

3.2 Solving the optimal hiring policy

In order to solve a firm's optimal hiring policy we need to set up its value function. Denote by V_0 the firm's value (expected present value of profits) in a period when it hires a novice manager. This is necessarily the case when it employed an experienced manager in the previous period. If it employed a novice in the previous period then it knows the ability θ of its incumbent manager and faces a choice between retaining him for another period or hiring a novice with unknown ability drawn from a known distribution.

There are three possible cases. If the incumbent manager is sufficiently bad, $\theta < c$, then by the assumed liquidity constraint he just caused the firm to go bankrupt and the value is zero. If the manager is not so disastrous but also not good enough to be rehired, $c \leq \theta < \Psi$, then the firm will discard him and hire a novice. In this case the value is by definition V_0. Finally, if the incumbent is good enough to be retained the firm will earn a profit θ - c this period and next period will revert to hiring a novice. The value of a firm with an incumbent manager of ability θ can therefore be written as,

$$\mathrm{V}(\theta) = \begin{cases} 0 & if\ \theta < c \\ \mathrm{V}_0 & if\ c \leq \theta < \psi \\ \theta - c + \delta V_0 & if\ \psi \leq \theta \end{cases} \qquad (1)$$

where δ in (0,1) is the discount factor.

To solve the firm's problem let's first eliminate the unknown ``initial" value V_0. Novices are drawn from a known distribution so V_0 is equal to the unconditional expectation $E[\pi(\theta) + \delta V(\theta)]$, where $\pi(\theta)$ is from (1):

$$V_0 = \Pr(\theta \geq \mathrm{c})\,(\mathrm{E}[\theta|\theta \geq \mathrm{c}] - \mathrm{c}) \\ + \delta\{\Pr(\psi > \theta \geq \mathrm{c})\,V_0 + \Pr(\theta \geq \psi)\,(\mathrm{E}[\theta|\theta \geq \psi] - \mathrm{c} + \delta \mathrm{V}_0)\} \qquad (2)$$

The last term takes into account that if the manager turns out to have $\theta > \psi$ then he is retained for another period and the firm will return to the ``initial value" V_0 two periods later.

In order to get a closed-form solution, let's assume that ability θ is distributed uniformly in [0,1]. (Also assume that c<0.5 to focus on the interesting case). Then expression (2) becomes

$$V_0 = (1-c)\left(\tfrac{1+c}{2} - c\right) + \delta\left\{(\psi - c)V_0 + (1-\psi)\left(\tfrac{1+\psi}{2} - c + \delta V_0\right)\right\}$$

$$\Rightarrow V_0 = \frac{(1-c)^2 + \delta(1-\psi)(1+\psi-2c)}{2(1 + \delta(c - \psi - \delta(1 - \psi)))} \qquad (3)$$

The optimal threshold ψ_c is obtained by maximizing expression (3) with respect to the retaining threshold ψ; the result is

$$\psi_c = \frac{1 - \delta(\delta - c) - \sqrt{1 - \delta(1-\delta^2)(1-2c+c^2) - \delta^2(1-c^2)}}{\delta(1-\delta)} \qquad (4)$$

At *c* = 0 the liquidity constraint is irrelevant and the firms' optimal rehiring threshold coincides with the efficient policy. By setting *c* = 0 in (4) we thus get (after simplification) the efficient threshold $\psi_0 = \left(1 + \delta - \sqrt{1+\delta}\right)/\delta$. This is clearly above the population mean 0.5 for any positive discount factor, so some above-average managers would not be rehired, and certainly no "substandard" managers would be rehired.

3.3 Discussion

Here we analyze how the firms' optimal hiring policy (4) depends on the severity of the liquidity constraint *c* and on the discount factor δ. Figure 1 plots the optimal rehiring threshold ψ_c as a function of *c* for selected values of δ. For example, using *c* = 0.1 and δ = 0.95 yields $\psi_c \approx 0.32$, so the firm would hold on to incumbents who are substantially worse than population average---even though novice managers, who are in expectation of average ability E[θ] = 0.5, could be hired at the same cost. The threat of bankruptcy causes firms to behave in a risk-averse manner and can cause them to hire known substandard managers in order to avoid the risk of disastrously bad managers.

The interaction of patience and liquidity constraint is not obvious, and is best understood by first considering the unconstrained special case. In the absence of a liquidity constraint *c* would just have been a fixed cost of operation without any impact on hiring policy. The efficient threshold $\psi_0 = \left(1 + \delta - \sqrt{1+\delta}\right)/\delta$ is increasing in patience δ and limits towards $2 - \sqrt{2} \approx 0.59$ as $\delta \to 1$. Hiring novices over mediocre incumbents is a type of an investment: less expected ability today, more option value tomorrow. A more patient planner would want to invest more, and this would require a stricter rehiring threshold.

When *c* > 0 then a different type of an "investment" incentive appears. The hiring of a known substandard manager leads to lower expected ability today, but yields a smaller probability of bankruptcy tomorrow. With higher *c* the fear of bankruptcy is more acute, and at moderate levels it leads firms to lower their standards in rehiring managers. However, the strength of this "fear motive" also depends on discounting: the more patient the owners the more averse they are to the risk of firm's value going to zero in the future.

The interaction of incentives and preferences for these two conflicting types of investments results in a non-monotonic relation whereby, for a given δ, the hiring threshold is initially made more lenient and eventually stricter as *c* grows. When neither discounting nor liquidity constraint are very high the privately optimal hiring policy is too lenient relative to social optimum, and even below the population mean for a wide range of parameters. However, with very low δ and high *c* the threshold is higher than is socially optimal, but this implies extremely high discount rates so is not likely to be realistic.

The most natural case here is arguably one where liquidity constrained firms find it in their interest to hire and retain experienced managers that are in expectation worse than novice managers. This turns out to be the empirically relevant case in our analysis of English football managers. This is also a case that could not be explained within the model of Terviö (2009), where the source of inefficiency is liquidity constrained workers and their inability to commit to long-term wage contracts.

4 Data and empirical methods

4.1 Dataset

In most activities, performance measurement is difficult and often imprecise, but in football, performance is a matter of official record. A manager’s career can stretch over hundreds and possibly thousands of games, yet there is no uncertainty about performance as measured by goals scored and conceded, which in turn determines the outcome of the game (loss, draw or win). Our data covers the goal difference for all games in the 38 seasons from 1973/74 to 2010/11 across the four professional English divisions, yielding more than 75,000 game observations in total.

Our data includes all identified club managers in our sample period, over 940 in total. Because of their significance to the success of the team, managers are closely scrutinized and their careers are well documented. Consequently, we were able to retrieve the manager’s identity for around 99% of all game observations. Table 1 summarizes three ways in which we can characterize football managers:

- Personal characteristics: age, experience (measured in games), experience in the English leagues, and nationality.
- The way they entered the managerial labor market. We measure:
  (a) whether they became a manager while still a player, i.e. a player-manager
  (b) whether they had job experience at a lower managerial level as an assistant manager or scout before their first management job,

(c) whether they were hired from within the club for their first spell,

(d) and in which division they entered.

- Their history as a player. We know:

  (a) if they played as a professional,

  (b) if they played in the four largest European leagues (top divisions of England, Germany, Italy and Spain),

  (c) the number of teams they played for in England,

  (d) if they played for the club they currently manage,

  (e) and if they represented their national team.

We present the sample mean and standard deviation for these variables, and the mean within quartiles based on total career length[3]. There is a very large variance in the experience levels of football managers. The mean experience at the end of a managerial career in our sample is 162 games over 36 months (almost 4 years given a 10 month playing season), but the top quartile is present for almost 450 games and over 100 months and the bottom quartile are present for only 10 games over 2 months. Most managers come from the British Isles, and foreign managers have on average shorter careers in our dataset. A significant fraction of managers began their managerial career while still playing for their club (player-managers), and this fraction increases with experience; in addition around one third of our sample was promoted to manager from within the club organization (but were not player-managers). The average division at the beginning of their careers is around 2.8 (out of 4), so that most managers "start at the bottom". About 95% of managers are former professional football players, two thirds played in one the big four European leagues and in their careers they played for between 3 and 4 clubs in England. Between one third and one half were managers at clubs they had once played for, while just over one third had played for their national representative team (suggesting they were among the most talented players).

Table 2 summarizes the data on the clubs, which employ the managers. Our financial data cover about 85% of all club-years during the sample years. We show both sample means and figures by quartiles in the average end-of-season league rankings. The first panel consists of financial data, which is highly skewed towards clubs in top quartile (i.e. mostly playing the top division, currently known as the Premier League). At the top end, revenues, wages and assets are all much larger, and the incidence of

[3] We provide detailed variable definitions in appendix B.

insolvency is lower, even though both at the top and bottom end clubs on average report pre-tax losses. In terms of sporting results, there are no significant inter-quartile differences but this is because competition is mostly intra-quartile (the quartiles correspond closely to the divisions). For example, a win percentage of 50% represents a much higher level of performance in the Premier League than in the fourth tier while the latter is a product, which is much less attractive to consumers. The bottom panel of Table 2 reports managerial hiring by teams. Tenure is slightly longer in the highest quartile. Clubs in the top quartile hire less frequently and are less likely to hire novice league managers, but are more likely to hire a foreign manager.

4.2 Rolling-window estimation

In our empirical analysis we identify the extent to which clubs rehire "substandard" experienced workers instead of hiring new (inexperienced) entrants. Unlike the stylized model above, manager careers last many periods and ability is not revealed at a single point in time. Firms learn about manager ability gradually over the manager's career. In our setting this process is explicit, as each game publicly reveals information on the ability of the managers involved. In our analysis we look at the ability of experienced managers who are rehired by clubs, and compare their revealed ability at the time of hiring to the ability of contemporary entrants, which the club could have hired in their place. This implies the need to estimate each rehired manager's ability, based solely on information revealed before the firm makes its hiring decision.

To achieve this, we divide the dataset in 380 months, labeled p, running from August 1973 to May 2011. We do not include June and July, as during these months no official league games are played. Our algorithm then runs through the dataset from month 100 (i.e. May 1983)[4] to 380 (May 2011), to estimate the ability of each incumbent manager at the end of each month. In other words, at the end of May 1983, we estimate the ability of each manager who had entered before the end of May 1983, taking only games played before this date into account. At the end of the next month (August 1983), we re-estimate manager ability taking the additional 4 or 5 games played in August into consideration. We repeat the process for September 1983, and so on until May 2011.

Through this "rolling window" estimation algorithm we obtain a profile of manager ability estimates over the manager's career, rather than a single, end-of-career estimate. For our analysis of hiring, we look at the ability estimate at the end of the last month before we observe the hire. As manager ability

[4] We use the first 10 seasons in the data as a learning period and start the analysis of hiring from August 1983.

is only revealed during employment, we focus our attention on instances where we observe a firm hiring a manager, who has at least had one previous employment spell in the data. We dub these managers 'rehired' managers.

To proxy the ability of potential entrants we generate an empirical equivalent of the entrant ability distribution using the contemporaneous ability estimates of recent entrants at the date of re-hiring. For each month p we collect all contemporaneous ability estimates for managers who entered their first managerial employment spell in the fifty months (five years) leading up to month p.[5] Both managers active and inactive in period p are taken into account to avoid selecting on labour market survival. We only consider managers who start their careers in England, as these are more representative of potential entrants into the English football industry.[6] Note that by these definitions, a manager coming in after a managerial career abroad is neither 'rehired', nor part of the entrant distribution. A manager with foreign experience can only be part of our analysis after an initial spell in England.

4.3 Assessing manager ability

We now need to obtain credible estimates for $\hat{\theta}_{mp}$, i.e. the empirical equivalent of manager productivity for manager m, estimated on information available in month p. Our primary measure of manager ability is based on the estimation of individual effects, as pioneered by Abowd et al. (1999).[7] In our case, we estimate a model explaining the goal difference, y_{ijt}, in a game between clubs i and j in season t. We relate this to a vector of logged inputs for each club (X_{it} and X_{jt}), a set of firm effects (γ_i and γ_j) and individual manager effects (μ_{mi} and μ_{mj}). Our baseline model is given as:

$$y_{ijt} = \beta_d X_{it} - \beta_d X_{jt} + \gamma_i - \gamma_j + \mu_{mi} - \mu_{mj} + \varepsilon_{ijt}. \tag{5}$$

As indicated by the d subscript, we allow the coefficients on variable inputs to differ by the division in which the game is played. The first term in vector X_{it} is a dummy denoting the home and away team,

[5] We vary this time frame to assess the robustness of our findings.

[6] In principle the pool of potential entrants is global in scope, but we do not have complete data on the global pool. Managers who come to England from abroad tend to be at upper end of the ability distribution, but we cannot say how representative they are of a set of potential entrants, and relative to managers from the UK they are relatively few in number. Thus if anything our entrant pool may understate the expected ability of the true entrant pool, implying that there are even more substandard rehires than we identify.

[7] Abowd et al (1999) apply employer and employee fixed effects to the estimation of individual wage equations and develop an identification strategy based on moving employees. While we apply the same identification strategy, our focus on firm output measures is more closely related to the analysis of Bertrand and Schoar (2004).

which is relevant because of the persistence of home advantage in professional team sports. As in Peeters and Szymanski (2014) we control for the player inputs at the manager's disposal by means of the total wage bill paid out by the club in season t. We further include the book value of tangible fixed assets to proxy the value of the stadium and training grounds, which is the only significant tangible asset in the financial accounts of the clubs.

Our approach to estimate the manager effects (labeled 'ability' below) mimics the methodology of Abowd et al. (1999) by introducing manager and firm dummies into the estimation of equation (5). To separately identify club and manager effects we require clubs to be connected into a network by at least one moving manager before or in month p. We first determine which clubs are part of the largest network using the routine of Cornelissen (2008), and then drop all observations pertaining to clubs that are not. Here, both clubs in the game observation have to be part of the connected network. We are using each game in the dataset twice, once from each team's perspective. This avoids the need to impose linear restrictions on the manager (club) dummies to ensure that these are estimated exactly opposite when the manager (club) takes on the role of either team i or team j. We cluster all standard errors at the level of individual games to correct for the interdependence among observations we introduce by using each game twice (for a similar approach, see e.g. Duggan & Levitt, 2002). In a final step, we standardize the ability measure for each manager by subtracting the mean ability estimate in the manager population in month p and dividing the resulting value by its standard deviation. This step is necessary to ensure the ability measure is comparable across estimation months, because the reference worker in the connected network, i.e. the 'zero' worker effect, may change as the rolling estimation progresses.

We have considered a host of alternatives to the baseline ability measure we specify in equation (5). Most notably, we calculate a cruder measure of manager ability, which simply compares a manager's performance to his predecessors at the club in the same division, without controlling for input use. This encompasses the possibility that a part of managerial ability is the capacity to funnel resources to improvements in the playing squad. We also estimate further variations of the baseline model, which (a) include controls for team-manager match quality, (b) put more weight on recent, rather than historical performance, or (c) condition on working experience. In the interest of clarity, we present results using the baseline ability estimate from equation (5) in all further analyses. We refer to appendix A for more detailed information and results for robustness checks we conducted using the alternative measures.

5 Main Empirical Results

5.1 Manager ability

In Table 3 we show summary statistics of the manager ability estimates at the level of individual managers. The first columns report the full sample numbers, which are then split in quartiles by the distribution of end-of-sample ability estimate. Mean ability equals 0.234 with a standard deviation of 0.835. This may seem odd given our standardization procedure, which subtracts the per-period average from the raw estimates. However, the table reports career averages and managers with higher ability tend to be accumulate more monthly observations. This is consistent with employers using the information provided by the market to screen employees and retain those who are more able, and who therefore enjoy longer tenure.

To provide a better feel for our rolling estimation, Figure 2 compares the evolution of our ability measure during the first ten years (100 months) in the career of three example managers, i.e. Arsène Wenger (Arsenal), Steve McMahon (Swindon, then Blackpool) and Micky Adams (among others, Fulham, Brentford, Leicester City and Coventry City). The gaps in the profiles of McMahon and Adams indicate unemployment spells, during which no information became available to update their estimates. As a reference, the shaded area depicts the interquartile range around the median of the manager ability estimates for managers with the same experience, expressed in months since labor market entry.[8] Naturally, the early career estimates for our example managers are quite noisy, because these are based on just a few observations. When comparing the abilities of the three managers, Wenger clearly comes out on top, with his revealed ability continuously in the upper end of the distribution. This will not come as a surprise to interested football followers, because Wenger is widely considered as one of the most influential managers in England. Micky Adams' ability estimates are much closer to the median of the manager workforce, which may have helped him to return to a new managerial job after his (multiple) unemployment spells. By contrast, McMahon's estimates continuously linger in the lower quartile, even though he managed to secure a new appointment (at Blackpool) after getting fired from his first job (with Swindon Town).

[8] This reference is indicative, but necessarily imperfect, because the reference population against which the effects are estimated changes over time. While our procedure to standardize by subtracting the contemporary mean and dividing by the contemporary standard deviation increases the comparability, we are still comparing relative performance in different reference groups.

5.2 How many hires of experienced managers are substandard?

Now we turn to comparing the ability estimates of rehired experienced managers with the distribution of entrant abilities. We use two bases for this comparison[9]- (i) the mean ability of an entrant in all divisions over the previous five years (of which there are on average about one hundred) and (ii) the mean ability of an entrant who started in the same division as the rehired manager (of whom there are typically around twenty five). The latter comparison addresses the concern that clubs in the lower divisions may not have access to the same pool of high quality novices as clubs in the higher divisions. In that case, we would overestimate the extent of substandard hiring in the lower divisions, as lower division clubs would be confined to hiring a marginal entrant of lower quality than the overall average, and, conversely, understate the extent of substandard hiring in the top divisions.

In Table 4, we show the results of this comparison. We have 755 hires of experienced managers for which ability can be estimated at the time of hiring. Our preferred comparison, which looks at the average of all entrants, is that approximately one quarter of all rehired managers are "substandard", i.e. their ability estimate at the time of hiring is below the average entrant. Across divisions, we find that between 15% and 37% of all rehired managers are substandard and there is a tendency for this proportion to be higher in the lower divisions. When looking at the comparisons within division, these numbers increase further in the higher divisions, while remaining relatively constant in the lower divisions. In other words, accounting for the varying quality of the entrant pool by division, strengthens our findings. We caution to over interpret this finding however, as the number of entrants in some comparisons is relatively small.

Figure 3 illustrates the extent of substandard rehiring. The upper panel plots the distribution of the rehired experienced managers' abilities measured in expected goal difference relative to the expected goal difference generated by the average entrant. All experienced managers with expected goal difference below zero are by definition substandard. Not only is there a large number of substandard rehires, but also a significant fraction of rehires are far below the mean. By way of comparison, our model estimates indicate that home advantage is worth a goal difference increase of 0.53 goals per game on average. To gain a goal difference of 0.5 through wage spending, a team has to outspend its opponent by about 2 to 1, all else equal.

[9] See appendix A for results using alternative ability estimates and comparison groups.

5.3 Are substandard managers underperforming?

A potential concern with any estimates of managerial ability is that firms may observe ability signals that are not in the data. If these omitted variables are important indicators of future performance, we may be classifying hires as substandard who in retrospect turn out to be successful. Thus, we need to examine whether we can reject the hypothesis that hires, which we classify as substandard, turn out to be justified by their subsequently improved performance.

We present the results of this analysis in Table 5, where we report the transition probabilities of rehired managers conditional on their rating as substandard or "above standard" at the start of their employment spell. As with table 4, we report comparisons based on the average ability of all entrants and entrants in the same division. We find that the rating of a rehired manager is unlikely to change by the end of their employment spell, 72% of managers rated substandard at the start of a spell were rated substandard at the end, while 94% of those rated above standard were still rated above standard by the end of their spell. This suggests that clubs have little reason for optimism that substandard manager will get significantly better, and every reason to favor managers who are already rated above standard. This is consistent with our argument that clubs will likely hire a substandard manager not because they expect to see an improvement, but because they are scared of the downside potential of an entrant.

5.4 When do firms hire substandard managers?

Next, we ask which firms hire substandard managers. To answer this question, we use the classification of firms in divisions as a proxy for both the productivity distribution and the severity of liquidity constraints. First, it is obvious that the output of each division (defined as games played) is approximately the same, but the interest of consumers and, consequently, the revenues generated per win are significantly larger in the higher divisions. In this sense, the *revenue* productivity is much higher in the higher divisions. Moreover, the divisions are connected through a system of promotion and relegation, which allows the best performing firms to move up to higher levels, while demoting poor performing firms to lower divisions. This implies that firms, which receive a positive shock, will improve their position according to our classification over time. Second, the extent of loss making and insolvency are more pronounced in the lower divisions. All but one of the 67 insolvency events in our sample period took place at clubs competing in the three lower divisions (See Szymanski (2015), chapter 8). Hence, clubs in the lower divisions are, on average, more likely to be liquidity constrained. Both these considerations would imply that, according to our theoretical framework, lower division clubs should be more inclined to hire substandard experienced managers.

We take a closer look at the firms and workers involved in substandard hiring through a simple regression model. As our unit of observation, we take each rehiring of an incumbent manager (m) by a firm (i) in a month (p), provided that a manager ability estimate is available for month p. Our dependent variable is an indicator, which takes value 1 if the incumbent is substandard in terms of his fixed effect estimate relative to the mean ability of an entrant over the previous five years, and 0 otherwise. We relate this variable to a vector of contemporaneous manager (X_{mp}) and firm characteristics (Z_{ip}) in a linear probability model. Our specification hence takes the form,

$$y_{mip} = \beta_x X_{mp} + \beta_z Z_{ip} + \varepsilon_{mip}. \quad (7)$$

For the manager characteristics, we first measure how many times firm i has been able to observe manager m on the job before month p. In column 1 of Table 6, we insert the log number of games the manager has coached in his career up until p, and how many of those games were played in England. We find that the smaller the number of observations on which a substandard manager's ability estimate is based, the more likely he is to be hired again. In column 2 we divide manager experience into bins to allow for non-linear effects, but find similar results. Managers with 40 to 80 games of experience (i.e. roughly between 1 and 2 seasons) are 18% less likely to be rehired as a substandard manager. When managers have been on the job for 80 games or more (beyond 2 seasons) this probability shrinks by 40% relative to the inexperienced reference category. These results indicate the importance of learning. Clubs are less willing to hire experienced substandard managers, which they have been able to observe extensively on-the-job.

In the last three columns we add indicator variables for the division in which the experienced manager was hired. These show that substandard managers are significantly more likely to be hired in the lower two divisions. This supports the notion that low productivity, liquidity constrained firms are more likely to hire substandard managers. Note that these clubs also tend to hire more entrants, as these are the firms facing a choice between experimenting and rehiring from the bottom end of the incumbent ability distribution.[10]

We included further covariates in the regression to test the robustness of these findings. We control for manager characteristics, such as age and nationality and introduce a measure for the month in which the hire took place to allow for trends in hiring (column 4-6). We then add variables to gauge

[10] We experimented with some indicators of financial distress to explore the relationship to substandard hires in more detail, but were unable to uncover any significant relationship. There are two primary reasons for this. First, conventional financial measures used to predict insolvency in other industries (e.g. profitability and asset/liability ratios) do not discriminate among clubs, since almost all clubs report losses and have weak balance sheets (see Szymanski (2017) for details). Second, the accounting information relating to clubs that effectively become insolvent is often missing, as clubs do not file accounts during insolvency procedures.

labor market entry mode (column 5), and previous playing career (column 6).[11] Most notably, we find that substandard rehiring is declining over time, that substandard rehires tend to be older and that substandard natives have more chance of getting rehired than substandard migrants.

5.5 What is the expected cost of hiring a substandard manager?

While we assume that firms act to maximize their self-interest, our reasoning still implies they pay a financial penalty, in the form of foregone revenues, for hiring substandard managers. Indeed, under certain circumstances firms choose to bear this cost to safeguard themselves from potential bankruptcy. This begs the question, however, how large the expected private cost of substandard hiring is.

To this end, we first quantify the effect of hiring a substandard manager on the club's output, i.e. its sports results defined as the end-of-season ranking. Based on our estimates of equation (5), we first calculate the difference in expected goal margin per game between the mean ability of all substandard managers and the mean ability of contemporary entrants. We then use historical results to translate this into the expected difference in end-of-season ranking. We find the mean entrant outperforms the mean substandard manager by about two thirds of a rank position e.g. from rank 10 to rank 10.66. While this may seem small, the number of teams in each division ranges from 20 to 24, so two thirds of a rank is between 3.3% and 2.75% of the range.

Next, we estimate a revenue function, which connects sports results to annual revenues, based on the specification of Peeters and Szymanski (2014).[12] By imputing our estimate for the difference in sports production between substandard and entrant managers in this function, we obtain Figure 4. In the upper panel, we show the hypothetical revenue impact from replacing a substandard manager by the average entrant for the median club in each division. Because of the extraordinary revenue growth football has experienced over the last 30 years, the financial loss of substandard hiring, in terms of foregone revenues, has also risen dramatically. For example, the average division 1 club would expect to lose about £0.5 million per year in the early eighties, versus about £8 million in 2011. Therefore, the lower panel depicts the same impact measured as a percentage of average annual revenues. While the estimated cost in relative terms has also risen in the top division, from about 7% to 9% of annual revenues, it has declined in the lower divisions to just over 1% in the second division and just under

[11] See appendix B for exact definitions of these variables.

[12] This specification relates revenues to sports results, tangible assets, firm and division-year fixed effects. Parameter estimates can be obtained from the authors.

1% in the lowest 2 divisions. Figure 4 clearly demonstrates that the foregone revenues resulting from a substandard hire are (a) non-trivial at all levels and (b) considerably larger for the top division, especially in later years. This also explains our finding above that top division clubs are more reluctant to hire substandard managers. Not only do they have a better supply of experienced managers to hire from, they also face greater financial incentives to avoid substandard managers.

5.6 How much less talented is the active workforce because of substandard hiring?

A central result of our model is that substandard hiring lowers the talent level in the population of active football managers. To measure the extent of this aggregate inefficiency, we construct a distribution of managerial ability following the counterfactual policy rule of never re-hiring managers, who are revealed to be substandard. In practice, we consider all the estimates $\hat{\theta}_{mp}$ for all managers *m* in months *p* when they were actually managing a club, but truncate the careers of substandard managers to their employment spells before they were rehired as a substandard experienced worker. As such, we assume counterfactually that substandard managers were never re-hired.[13] We fill the resulting reduction in manager-months by adding counterfactual ability estimates, which we draw from the observed distribution of estimated abilities in the non-substandard entrant population. To aid interpretation, we scale the estimated abilities by the standard deviation in the population of estimated abilities. We then calculate the average ability in the actual and counterfactual workforce at the level of the individual manager-month, i.e. a manager's rolling ability estimates appear once in the population for each month he has been active in the actual or counterfactual population. The comparison between these averages is the central result of our counterfactual exercise.

Table 7 summarizes the results of this counterfactual. Out of the 22,697 manager-months for which we can estimate manager ability using the rolling window, we exclude 4,324 in the counterfactual, because they belong to the end-parts of substandard worker careers. In total, we replace the end-parts of 100 managerial careers by 271 counterfactual employment spells, which increases the total number of managers in the counterfactual population by 171. The average career length in the worker population drops by 11 about months, both because of the introduction of short-lived entrants and the curbing of incumbent careers. The average ability of the counterfactual managers is 0.37 standard

[13] For example, in Figure 2, McMahon is estimated to be substandard 50 months after entry and at the start of his second employment spell; therefore only his first spell is included in the counterfactual data of manager-months. His other spells are replaced by the average entrant ability in the year following the start of his employment spells.

deviations higher than the average ability of the substandard managers they replace. A bootstrapping exercise using 1000 replications of the entire procedure, suggests that this difference is indeed significant.[14] Looking at the full workforce, the difference drops to 0.07 standard deviations. This is because we replace a relatively small share of the total population of manager-months (19%, i.e. 81% of manager-months are exactly equal in both populations).

In Figure 5, we plot the distribution of ability estimates for the actual and counterfactual workforce, where the level of observation is an active manager-month in the dataset. In panel A, we focus on the subsample of employment spells we replace in the counterfactual. The graphs show how the counterfactual dismisses a large density of experienced managers of low to modest ability. The entrant population, which replaces them, has a wider variance in ability, but on average outperforms the experienced managers they replace. In other words, firms face more uncertainty about the ability of these managers pre-hiring, but, as shown in the right tale of the distribution, the counterfactual entrants possess upside potential, which is absent in the actual data. In panel B, we depict the same comparison for the full sample of observed manager-months in the data. These graphs illustrate that the counterfactual policy has a modest effect on the overall distribution of ability in the active labor force.

6 Supporting Results

6.1 Manager ability at entry and the speed of learning

As can be seen in Table 1, a good deal is known about most managers at the beginning of their career, even when managerial ability may not be observed. It seems reasonable to assume that the wealth of attributes teams (and econometricians) can readily observe, contains some information to help predict managerial ability. Moreover, for our model to make sense, future estimates of managerial ability has to be predictable after a manager has entered the labor market. We examine these assumptions by regressing the current estimate of managerial ability on the estimate of managerial ability (a) in the previous month and (b) five months earlier. Thus our forecasting model for entrant ability is,

$$\hat{q}_{mp+k} = \beta_x X_{mp} + \beta_q \hat{q}_{mp} + \varepsilon_{mp}. \tag{8}$$

[14] Some caution is required in interpreting these results, since we treat estimated abilities as given in the bootstrap replications. In other words, the estimation error induced on the ability estimates by the rolling window estimation is not fully taken into account. However, bootstrapping the entire rolling estimation procedure would take up a prohibitive amount of computer time.

Here, $\hat{q}$ is the estimate of ability and k is the number of time periods ahead (one or five) and the vector X_{mp} includes manager characteristics in month p, in particular, log manager age, its square, the log month of entry, a dummy for foreign managers and all variables in Table 1 which refer to the previous playing career and labor market entry mode. To proxy the predictability of ability estimates before labor market, we also run this regression excluding information we observe on the job, i.e.

$$\hat{q}_{m,p+k} = \beta_x X_{m,p} + \varepsilon_{mp}. \tag{9}$$

We focus on the adjusted R^2 of these regressions to identify the extent of learning. The graphs of the results for 1 and 5 months ahead are shown in Figure 6. Both graphs tell a strikingly similar story. In the one month ahead model the R^2 jumps from around zero at entry to about 0.8 in the first month and by ten months has risen to over 0.95 where it stabilizes. With the five month ahead model the R^2 jumps from around zero at entry to about 0.6 in the first month and by twenty months has risen to over 0.9 where it stabilizes. In both cases the R^2 of the model based on pre-entry characteristics never moves significantly away from zero. These results suggest that learning is relatively fast in this industry. By comparison, Lange (2007) reports that the initial expectation errors of employers fall by 50% within 3 years based on an annual survey.

An important caveat to this analysis is that we are restricted to assess learning from the point-of-view of the econometrician. Employers may have access to private ability signals, which we cannot observe. In our view, this implies that in reality firms could learn even faster than we estimate to be the case. In that sense, our assessment puts a lower bound on the speed of employer learning.

6.2 Manager ability and attrition

If employers are able to observe the ability of their own manager with increasing accuracy over time, we should expect them to terminate their manager's employment when he turns out to be of low ability. In Table 8, we look at this relationship by estimating a linear probability model of employment termination. The regression is of the form

$$y_{mip} = \beta_q \hat{q}_{mp} + \beta_{dq} \Delta\hat{q}_{mp-1} + \beta_x X_{mp} + \varepsilon_{mp}. \tag{10}$$

The dependent variable, y_{mip}, is an indicator, which equals 1 if the manager's tenure at team i ended in month p.[15] We explain spell termination by two main variables, (a) the current ability estimate of the manager, $\hat{q}_{mp}$, and (b) the update in the ability estimate, relative to the previous month, $\Delta\hat{q}_{mp-1}$. We further control for a vector of manager characteristics containing log of manager age, tenure and experience and add FEs for the year, calendar month and division.

Table 8 confirms that managers with higher ability estimates have a significantly lower probability that their employment spell is terminated. An improvement by one standard deviation leads to a decrease in termination probability between 1.9% (for the base) and 2.6% (for the add win%), while the baseline probability is around 5%. A positive update in the ability estimate relative to month $p-1$ adds further to this effect. These results also appear very robust to adding the other personal characteristics to the model. In conclusion, clubs are more likely to retain high ability than low ability managers, and even more likely to retain them when they are on an upward trajectory.

6.3 Manager ability and career progress

Finally, we assess the observability of manager ability to rival firms. First, we examine whether higher estimated managerial ability increases the probability of being retained in the labor market of football managers. We estimate a linear probability model where the dependent, y_{mp}, indicates that manager m is rehired by any other team in any country after the end of his spell in month p. We insert the estimated manager ability in month p, and a vector of personal characteristics, X_{mp}, as explanatory variables, i.e.

$$y_{mp} = \beta_q \hat{q}_{mp} + \beta_x X_{mp} + \varepsilon_{mp}. \tag{11}$$

Second, we estimate an ordered probit model of career progress using the same set of explanatory variables. Here we model the probability of four categorical outcomes in terms of career progress, i.e. not being rehired at all, being rehired in a lower division than the one you previously worked in, being rehired in the same division, and being rehired into higher division.

Table 9 reports the model estimates for our sample. The model variants display a strikingly consistent pattern: managers with higher ability estimates on ending an employment spell are more likely to be rehired and more likely to be rehired in a higher division. This supports the notion that clubs act on the

[15] In an attempt to separate voluntary 'quits' from involuntary 'firings', we also estimate models where the dependent variable only equals 1 if the manager is not subsequently hired by a club playing in a higher division than his current club. This yielded equivalent results, which are available on request.

information revealed over time, both when rehiring incumbents and when poaching workers from lower productivity clubs.

We also find that rehires and career progression are more likely for more experienced managers, younger managers and managers that are foreign. This last observation may be due to the fact foreign managers would usually need to be better than a domestic alternative when first being hired, either for objective reasons (e.g. language skills) or prejudice, while over time those disadvantages are likely to diminish (while ability remains stable).

7 Discussion and conclusions

The core finding in this paper is that firms in labor markets, where worker ability can only be observed "on-the-job", have an inefficient incentive to hire experienced workers of low ability instead of labor market entrants.

The previous literature on hiring with on-the-job learning has assumed that firms face no liquidity constraints but have to fear other firms "poaching" the best talents they uncover. In Terviö (2009) firms rehire any experienced manager who is expected to be at least as able as a new manager; this is rational for individual firms but socially inefficient as untried managers have the potential to turn out much better than the true and tried mediocrities. However, the fear of poaching is not sufficient to explain why worse-than-average incumbents are hired over entrants with *higher* expected ability. In this paper we presented a model that shows why the presence of liquidity constrained firms can tilt the hiring of managers even more in favor of incumbents than the fear of poaching---and eventually to the detriment of the average quality of managers in the industry.

In our empirical analysis, we find support for this inefficiency in the English labor market for football managers. We exploit a very rich dataset on this labor market, covering complete employment histories, worker characteristics, firm financials and firm output (i.e. game results) over a period of 38 years. The quality of the data allows us to introduce an explicit model of employer learning on worker ability. Since we know the exact date and result of each game - the ultimate yardstick of managerial performance - we can infer which information on worker performance "on-the-job" was available to employers at each point in time. We empirically exploit this feature in a novel algorithm, which updates estimates of worker effects à la Abowd et al (1999) after each calendar month in the dataset, taking newly revealed observations on manager performance into account. As such, we generate a career profile for each manager's estimated ability over time, rather than a single estimate based on his full career. We then investigate each instance where a firm hires an experienced manager, and compare the current ability estimate of this experienced manager to the average current ability estimates of recent labor market entrants. In this way, we avoid hindsight in our analysis of firm hiring behavior.

Our results indicate that for about one quarter of all instances where a firm hires an experienced manager, the experienced manager has an estimated ability below the average entrant at the time of hiring, a situation we refer to as a 'substandard' hire. We show that these substandard hires are more common in lower divisions, where firms are both more liquidity constrained and the foregone revenues resulting from a poor hire are substantially lower. This observation is therefore in line with both our own theoretical predictions and those of Terviö (2009). In a counterfactual exercise, which replaces the substandard managers in our sample by hypothetical new entrants, we find that this policy would lead to a significant increase in the average worker ability of the active workforce over our sample period. As such, this paper provides the first empirical evidence of inefficient hiring using observational data from a high profile labor market. We thus add to the evidence reported by Pallais (2014), who documents that firms hire inefficiently little labor market entrants through a field experiment in an online job portal.

We finally generate a set of additional results, which support the underlying assumptions of the theoretical framework we use in our analysis. We show that worker ability is indeed very hard to predict before managers enter the labor market, but quickly revealed during their career. Then, we test the notion that the managerial labor market functions as one would expect if it were populated by maximizing agents. For example, managers with higher estimated ability have longer employment spells, are more likely to be rehired after an employment spell ends, and more likely to be hired by clubs operating at higher levels of competition.

The relative efficiency of labor markets is an important policy issue, given the increasing role of human capital in modern economies. This suggests that it is important not only to identify inefficiencies but also to measure their size. In this paper, we have effectively assumed that the relative ability of managers in winning games is also a measure of their relative ability in producing economic value. If the whole industry were only engaged in zero-sum competition, as it clearly is in sporting terms, then hiring decisions would not have any implications for aggregate efficiency. However, we think it is reasonable to assume that manager ability is also associated with contributions to economic value, because total revenue and total consumer surplus generated by the football league system are not fixed. Better managers (just like better players) contribute to higher quality of the sporting competition, which in turn affects revenue and consumer surplus. In this sense, our findings suggest the existence of a market failure, in that a significant fraction of rehires are inferior to the mean entrant. In light of our model, a tendency to hire managers that have shown themselves to be below average but not catastrophically bad is optimal from individual clubs' perspective, but in the longer run it leads to lower average quality of football managers.

The fact that football management is a highly specific task raises the problem of generalizability; however, this is likely to be the case for almost any activity involving high-skilled workers. As has long been recognized (e.g., Kahn, 2000), sports markets offer an outstanding laboratory for research on economic issues. The research in this paper is a clear example of this, as it exploits the richness of data available in the sports labor market to push beyond what is empirically feasible in other settings.

References

Abowd, J. M., Kramarz, F. & Margolis, D. M., 1999. High Wage Workers and High Wage Firms. *Econometrica,* 67(2), pp. 251-333.

Audas, R., Dobson, S. & Goddard, J., 1999. Organizational performance and managerial turnover. *Managerial and Decision Economics,* 20(6), pp. 305-318.

Bachan, R., Reilly, B. & Witt, R., 2008. The Hazard of Being an English Football League Manager: Empirical Estimates for Three Recent League Seasons. *The Journal of the Operational Research Society,* 59(7), pp. 884-891.

Barros, C. P., Frick, B. & Passos, J., 2009. Coaching for Survival: The Hazards of Head Coach Careers in the German Bundesliga. *Applied Economics,* 41(25), pp. 3303-3311.

Beech, J., Horsman, S. J. & Magraw, J., 2010. Insolency events among English football clubs. *International Journal of Sports Marketing & Sponsorship,* 11(3), pp. 236-249.

Bell, A., Brooks, C. & Markham, T., 2013. The Performance of Football Club Managers: skill or luck?. *Economics and Finance Research,* 1(1), pp. 19-30.

Bertrand, M. & Schoar, A., 2003. Managing with Style: the Effect of Managers on Firm Policies. *Quarterly Journal of Economics,* CXVIII(4), pp. 1169-1208.

Bridgewater, S., Kahn, L. M. & Goodall, A. H., 2011. Substitution and complementarity between managers and subordinates: Evidence from British football. *Labour Economics,* Volume 18, pp. 275-286.

Bruinshoofd, A. & ter Weel, B., 2003. Manager to Go? Performance Dips Reconsidered with Evidence from Dutch Football. *European Journal of Operational Research,* 148(2), pp. 233-246.

Cornelissen, T., 2008. The Stata command felsdvreg to fit a linear model with two high–dimensional fixed effects. *Stata Journal,* 8(2), pp. 170-189.

d'Addona, S. & Kind, A., 2014. Forced Manager Turnovers in English Soccer Leagues: A Long-Term Perspective. *Journal of Sports Economics,* 15(1), p. 107–131.

De Paola, M. & Scoppa, V., 2011. The effects of managerial turnover: evidence from coach dismissals in Italian soccer teams. *Journal of Sports Economics,* 13(2), pp. 152-168.

Duggan, M. & Levitt, S. D., 2002. Winning Isn't Everything: Corruption in Sumo Wrestling. *American Economic Review,* 92(5), pp. 1594-1605.

Fama, E. F., 1980. Agency Problems and the Theory of the Firm. *Journal of Political Economy,* 88(2), pp. 288-307.

Frick, B., Barros, J. P. & Prinz, J., 2010. Analysing Head Coach Dismissals in the German Bundesliga with a Mixed Logit Approach. *European Journal of Operational Research,* 200(1), pp. 151-159.

Frick, B. & Simmons, R., 2008. The impact of managerial quality on organizational performance: evidence from German soccer. *Managerial and Decision Economics,* 7(593-600), p. 29.

Goodall, A. H., Kahn, L. M. & Oswald, A. J., 2011. Why do leaders matter? A study of expert knowledge in a superstar setting. *Journal of Economic Behavior & Organization,* Volume 77, pp. 265-284.

Jackson, K. C., 2013. Match Quality, Worker Productivity, and Worker Mobility: Direct Evidence from Teachers. *The Review of Economics and Statistics,* 95(4), p. 1096–1116.

Jovanovic, B., 1979. Job matching and the theory of turnover. *Journal of Political Economy,* 87(5), pp. 972-990.

Kahn, L. B., 2013. Asymmetric Information between Employers. *American Economic Journal: Applied Economics,* 5(4), pp. 165-205.

Kahn, L. B. & Lange, F., 2014. Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures. *Review of Economic Studies,* Volume 81, p. 1575–1613.

Kahn, L. M., 2000. The Sports Business as a Labor Market Laboratory. *Journal of Economic Perspectives,* 14(3), pp. 75-94.

Koning, R., 2003. An Econometric Evaluation of the Effect of Firing a Coach on Team Performance. *Applied Economics,* Volume 35, pp. 555-564.

Lange, F., 2007. The Speed of Employer Learning. *Journal of Labor Economics,* 25(1), pp. 1-35.

Lazear, E. P., Shaw, K. L. & Stanton, C., 2015. The Value of Bosses. *Journal of Labor Economics,* 33(4), pp. 823-861.

MacDonald, G. M., 1988. The economics of rising stars. *American Economic Review ,* 78(1), pp. 155-166.

Muehlheusser, G., Schneemann, S., Sliwka, D. & Wallmeier, N., 2016 . The Contribution of Managers to Organizational Success: Evidence from German Soccer. *Journal of Sports Economics forthcoming.*

Pallais, A., 2014. Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review,* 104(11), p. 3565–3599.

Peeters, T., Salaga, S. & Juravich, M., 2015. *Matching and Winning? The Impact of Upper- and Middle Managers on Team Performance,* Rotterdam: Tinbergen Institute Discussion Paper.

Peeters, T. & Szymanski, S., 2014. Financial Fair Play in European Football. *Economic Policy,* 29(78), pp. 343-360.

Schneemann, S., Muehlheusser, G. & Sliwka, D., 2016. The Impact of Managerial Change on Performance. The Role of Team Heterogeneity. *Economic Inquiry,* 54(2), pp. 1128-1149.

Schönberg, U., 2007. Testing for Asymmetric Employer Learning. *Journal of Labor Economics,* 25(4), pp. 651-691.

Stanton, C. & Thomas, C., 2016. Landing the First Job: The Value of Intermediaries in Online Hiring. *Review of Economic Studies,* 83(2), pp. 810-854.

Szymanski, S., 2015. *Money and Soccer: a soccernomics guide.* 1 ed. New York: Nation Books.

Szymanski, S., 2017. Entry into exit: insolvency in English professional football. *Scottish Journal of Political Economy Forthcoming.*

Tena, J. d. D. & Forrest, D., 2007. Within-season Dismissal of Football Coaches: Statistical Analysis of Causes and Consequences.. *European Journal of Operational Research,* Volume 181, pp. 362-373.

Terviö, M., 2006. Transfer Fee Regulations and Player Development. *Journal of the European Economic Association,* 4(5), pp. 957-987.

Terviö, M., 2009. Superstars and Mediocrities: Market Failure in the Discovery of Talent. *Review of Economic Studies,* Volume 76, p. 829–850.

Van Ours, J. C. & Van Tuijl, M. A., 2016. In-season Head-Coach Dismissals and the Performance of Professional Football Teams. *Economic Inquiry,* 54(1), pp. 591-604.

Tables and Figures

Figure 1: Firm optimal retaining threshold ψ_c as a function of liquidity constraint c at various levels of discount factor δ, when entrant ability is uniformly distributed on [0,1]. Cases below the dashed line (population mean) indicate substandard hiring.

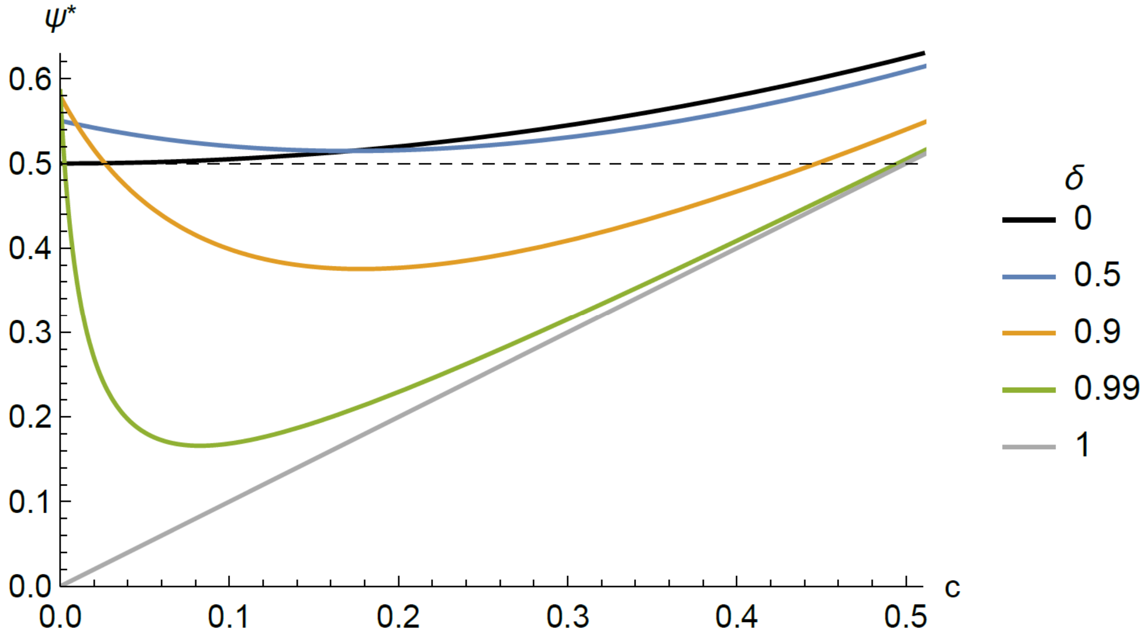

Figure 2: Career evolution of ability estimates for selected managers

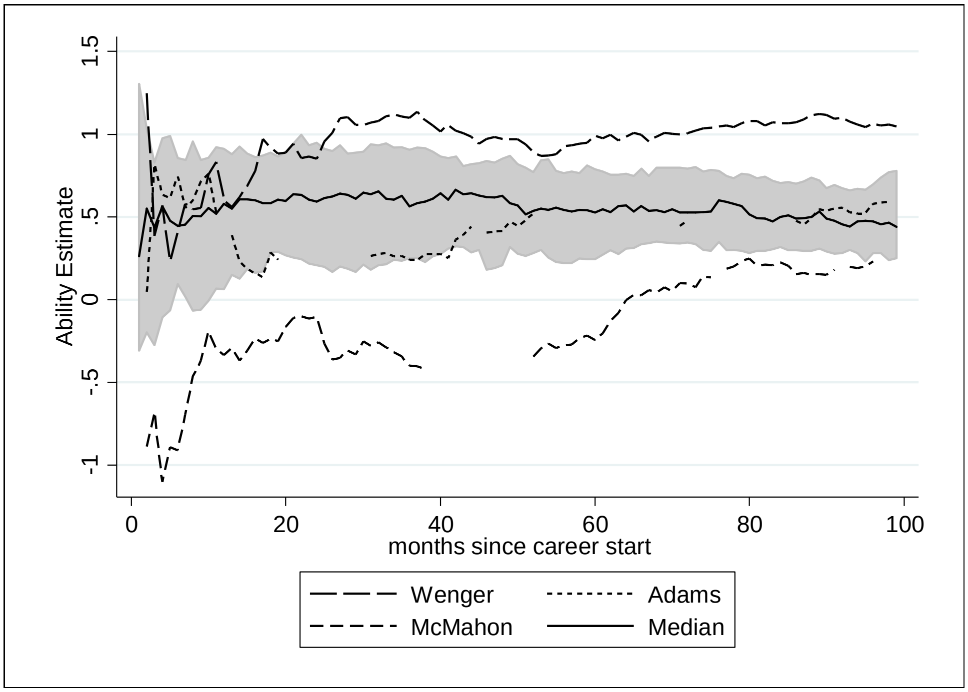

<u>Notes:</u> The plots depict the evolution of the ability estimates of Arsène Wenger, Micky Adams and Steve McMahon over the first ten years of their careers. The shaded area is the interquartile range for the ability estimates of managers with the same experience in number of months since labor market entry.

Figure 3: Ability estimates of rehired experienced managers versus average of recent entrants

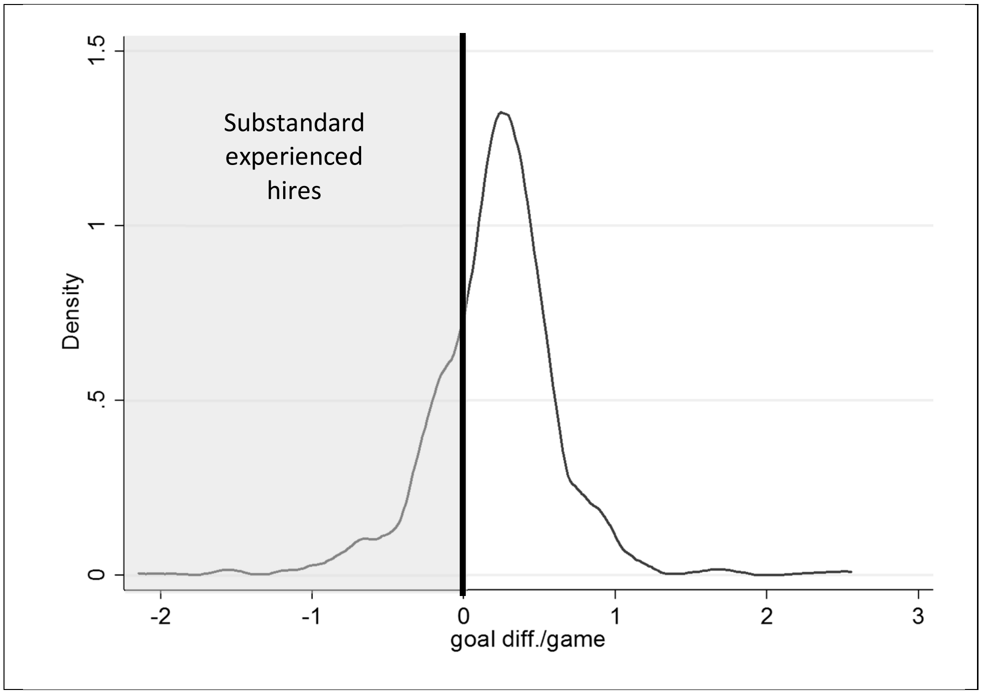

Notes: The figure depicts the distribution of the ability estimates of rehired managers vis-à-vis the average entrant. We graph the Kernel density for the estimated worker ability minus the mean entrant ability at the time of hiring, expressed in terms of goal difference produced per game. We show results using a 5 year time window of past entrants, other ability measures and choices for the entrant distribution can be provided on request.

Figure 4: Revenue impact of hiring average entrant vs. substandard experienced manager

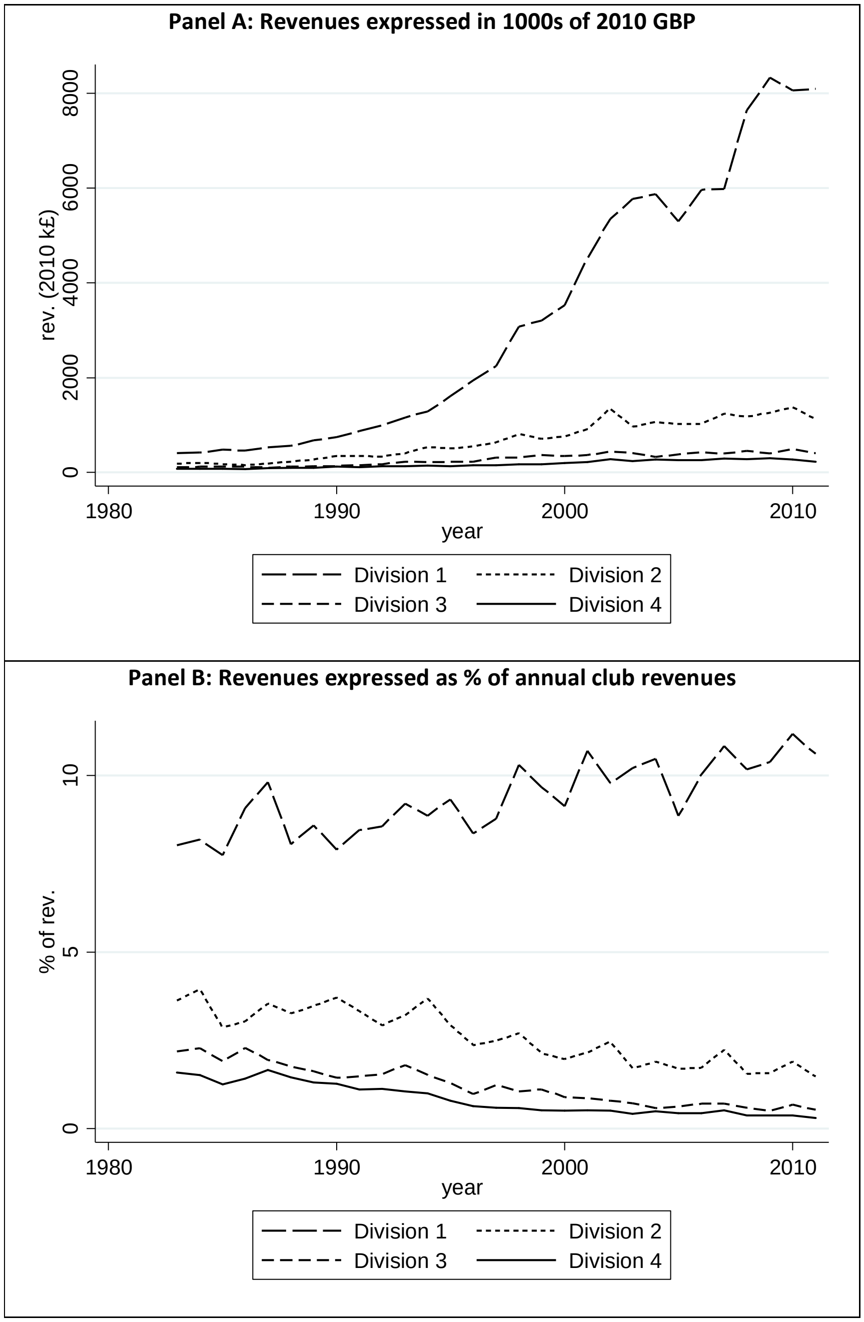

Notes: The graphs show the revenue impact of hiring the average entrant over the past 5 years instead of the substandard experienced managers, who were actually hired in the data. Figures are based on revenue regressions as in Peeters and Szymanski (2014), where revenues depend on tangible assets, goal difference, relegation or promotion in previous year, having a parent company and year-division as well as club effects. Graphs show the effect for a representative club in each division and year, which is constructed by taking the median in terms of asset holdings, club fixed effect and goal difference, and assuming the club has not been promoted or relegated and has no parent company. The top panel reports the absolute revenue effect in 1000s of 2010 GBP, the bottom panel looks at the effect as a fraction of annual club revenues.

Figure 5: Distribution of manager ability estimates in actual and counterfactual workforce

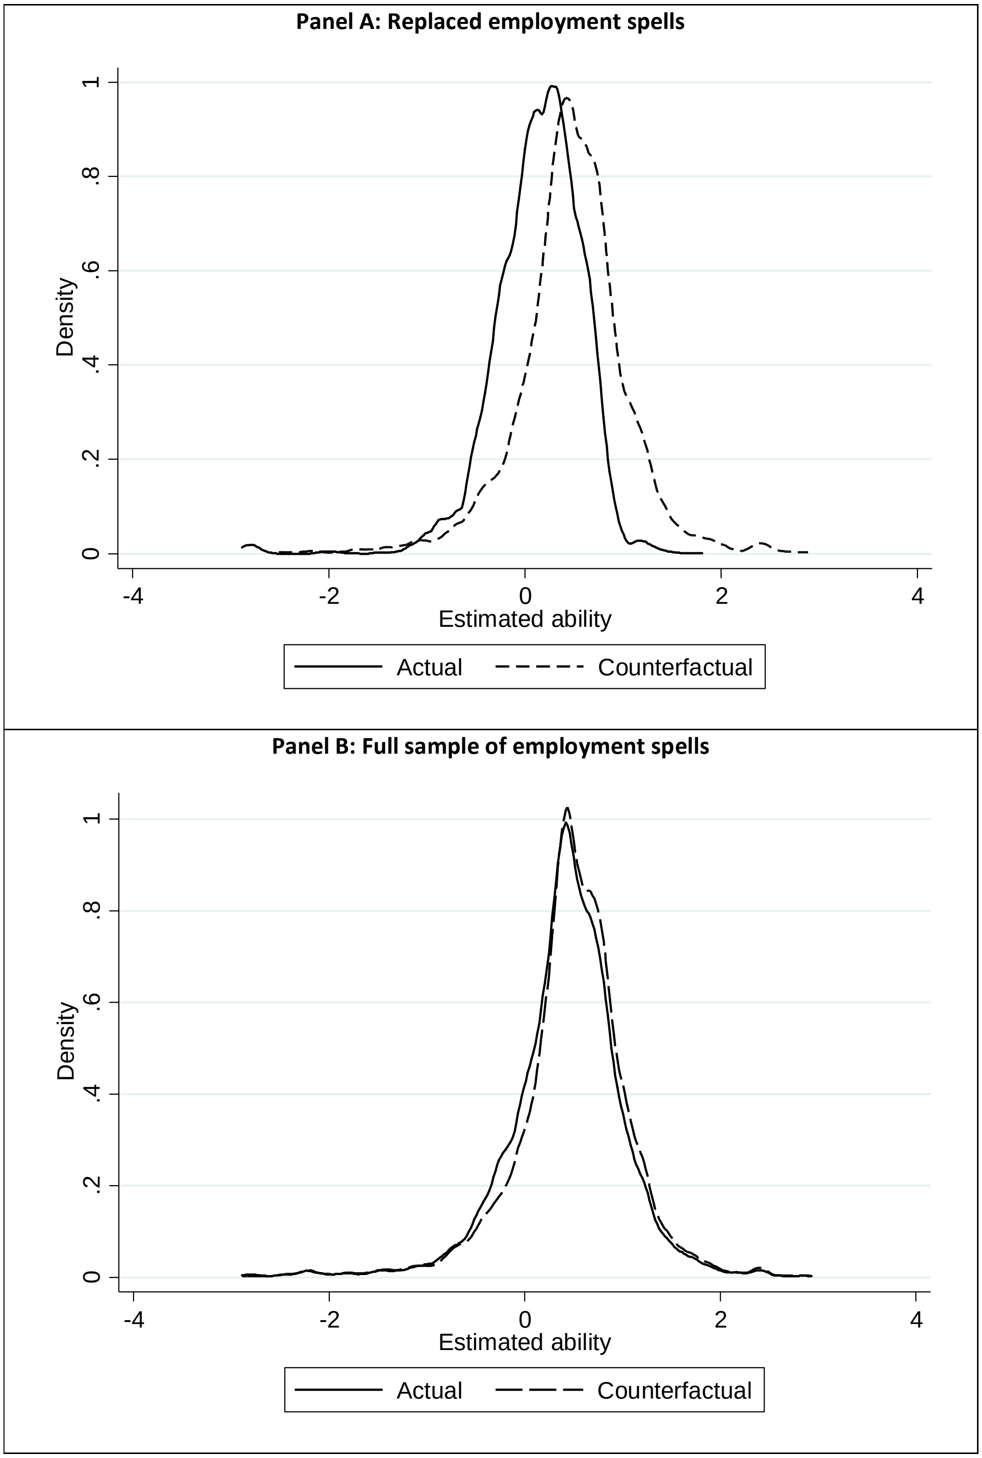

Notes: The graphs depict the distribution of contemporary ability estimates in the actual and counterfactual worker populations. The top panel focuses on the replaced subsample, i.e. end career spells of substandard managers. The bottom panel depicts the full sample results. The level of observation is a manager-month, so each manager appears once for each month he has been active in the sample. Ability estimates are expressed in standard deviations of the full sample of ability estimates.

Figure 6: Adjusted R-squared for regression of manager ability estimates 1 and 5 months ahead on entry characteristics and current ability estimates, by number of months observed

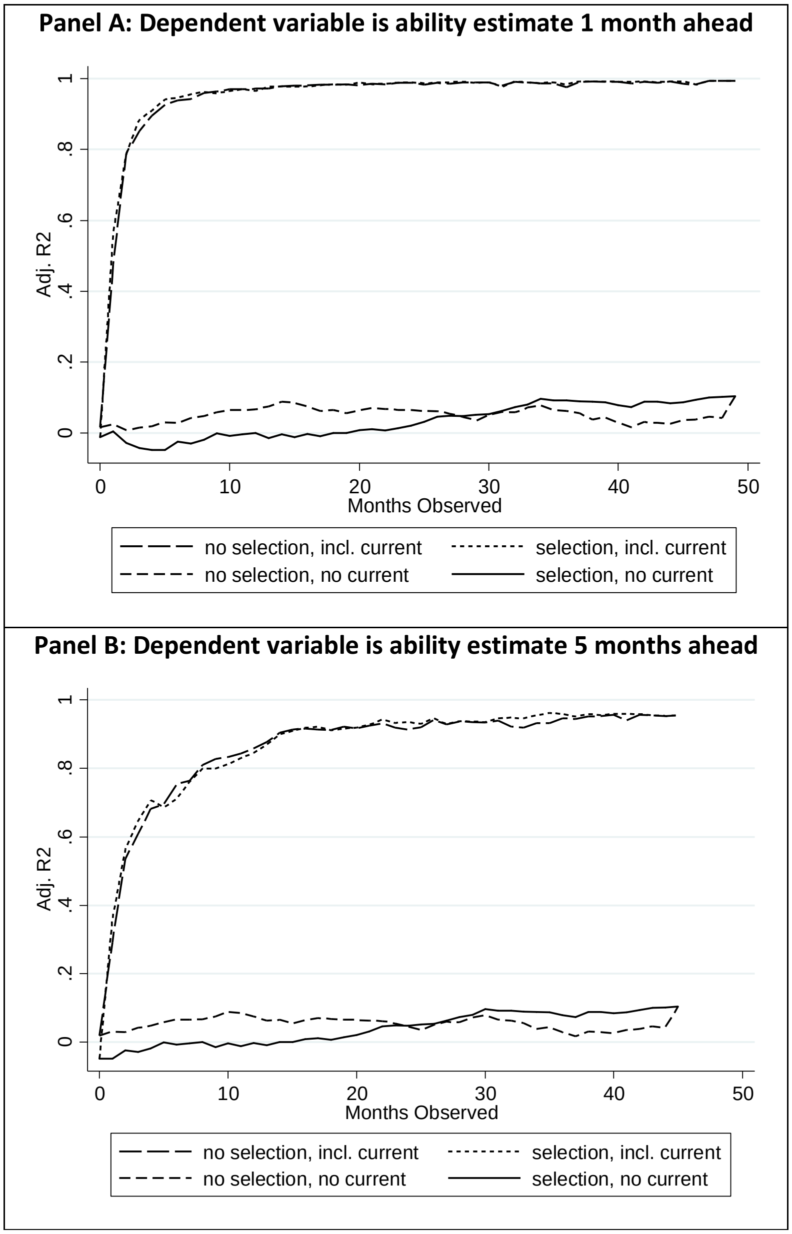

__**Notes**__**:** To form these graphs we draw subsamples of the dataset by the number of months a manager has been observed in the data. In each subsample we regress worker ability estimates 1 month ahead (top panel) and 5 months ahead (lower panel) on (a) a set of worker characteristics at labor market entry (label: 'no current') and (b) the same set of entry characteristics plus the worker's current ability estimate (label: 'incl. current'). The figures compare the adjusted R-squareds obtained by these models in each subsample, both for all managers ('no selection') and for the subset of managers that attains at least 50 month observations ('selection').

Table 1: Summary statistics at individual manager level by quartile in number of games present in dataset

	Full sample			Obs. quart 1		Obs. quart 2		Obs. quart 3		Obs. quart 4	
Variables	**Obs.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Observations											
Games observed	942	161.7	201.0	10.4	7.3	51.5	16.3	138.8	40.3	448.2	206.0
Months observed	942	36.4	46.0	2.3	1.5	11.3	3.8	30.9	9.2	101.6	47.7
Manager Char.											
Av. Age (year)	819	43.0	6.51	43.0	7.75	41.8	7.09	42.3	6.29	44.7	4.89
Av. Exp. (game)	830	142.0	166.1	66.7	158.8	79.8	153.1	115.0	113.5	272.0	152.0
Av. Eng. Exp. (game)	830	116.0	141.2	27.2	98.8	49.1	101.3	88.3	76.8	257.0	139.8
Foreigner	830	0.060	0.240	0.098	0.300	0.078	0.270	0.060	0.240	0.021	0.140
Market Entry											
Player-manager	827	0.295	0.460	0.244	0.430	0.240	0.430	0.297	0.460	0.375	0.490
Other man. exp.	827	0.518	0.500	0.580	0.500	0.555	0.500	0.500	0.500	0.464	0.500
Intern hire	827	0.335	0.470	0.458	0.500	0.362	0.480	0.322	0.470	0.251	0.430
Division	942	2.786	0.980	2.807	1.050	2.978	1.000	2.815	0.970	2.543	0.850
Playing history											
Play prof.	830	0.954	0.210	0.910	0.290	0.952	0.210	0.961	0.190	0.975	0.160
Play big 4	830	0.665	0.470	0.602	0.490	0.683	0.470	0.625	0.490	0.723	0.450
Num. Eng. Team	830	3.460	2.292	3.290	2.631	3.620	2.353	3.360	2.298	3.510	2.004
Ex-player club	830	0.440	0.448	0.530	0.499	0.460	0.489	0.460	0.459	0.330	0.330
International	830	0.375	0.480	0.308	0.460	0.339	0.470	0.431	0.500	0.392	0.490

Notes: Table depicts summary statistics for all manager characteristics in the dataset at the level of the individual manager. We provide statistics for the full sample and a breakdown by quartiles in the number of games we observe the manager.

Table 2: Summary statistics club-level data in estimation sample

Subsample:	Full sample			Rank quart 1			Rank quart 2			Rank quart 3			Rank quart 4		
Variables	**# Clubs**	**Mean**	**St. Dev.**	**# Clubs**	**Mean**	**St. Dev.**	**# Clubs**	**Mean**	**St. Dev.**	**# Clubs**	**Mean**	**St. Dev.**	**# Clubs**	**Mean**	**St. Dev.**
Financial data (2010 k£)															
Years observed	98	28.8	8.83	24	20.8	9.53	25	29.3	8.89	24	31.1	6.24	25	33.6	4.27
Wages	98	4731	6993	24	503	454	25	1235	720	24	4120	2763	25	12874	9463
Revenue	98	7304	12336	24	701	711	25	1767	936	24	5428	3848	25	20981	18042
Fixed assets	98	10757	23606	24	657	645	25	1448	967	24	5948	4812	25	34380	37842
Pos. net assets	98	58.6%	28.8%	24	53.8%	31.6%	25	56.6%	28.9%	24	47.2%	27.1%	25	76.0%	19.5%
Pre-tax profits	98	-970	2851	24	-26.5	126	25	-229	390	24	-990	1083	25	-2598	5228
# insolvencies/club	98	0.67	0.87	24	0.75	0.85	25	0.88	1.05	24	0.63	0.88	25	0.44	0.65
Sports results															
Goals pro	98	57.8	3.59	24	58.0	3.49	25	59.2	2.60	24	57.1	2.98	25	57.1	4.71
Goals against	98	57.5	5.75	24	62.9	3.55	25	59.0	1.94	24	57.2	2.88	25	51.3	6.26
Goal Difference	98	0.30	7.22	24	-4.89	5.31	25	0.23	3.10	24	-0.12	3.28	25	5.76	10.27
Points	98	57.3	4.70	24	54.9	4.76	25	58.9	3.96	24	57.3	3.56	25	57.9	5.54
Win %	98	50.2%	3.8%	24	47.8%	2.6%	25	50.1%	1.9%	24	49.9%	2.0%	25	53.0%	5.6%
Managerial hiring															
Number hires	98	16.9	6.88	24	12.9	6.34	25	18.2	6.89	24	18.9	5.94	25	17.7	6.97
Av. hires/year obs.	98	0.60	0.19	24	0.64	0.17	25	0.64	0.19	24	0.60	0.15	25	0.53	0.21
Av. manager tenure	98	182	96.3	24	177	90.7	25	163	73.6	24	165	48.2	25	222	140
# entrant hires	98	7.89	3.91	24	7.96	3.96	25	9.56	3.98	24	8.42	3.90	25	5.64	2.78
# foreign hires	98	0.61	1.15	24	0.21	0.41	25	0.40	0.76	24	0.54	0.72	25	1.28	1.86

Notes: Table shows summary statistics for club level variables in the dataset at the level of the individual club. We provide statistics for the full sample and a breakdown by quartiles in the average end-of-season ranking the club obtains over the sample period. All clubs are limited liability companies registered in the UK, which have to deposit a copy of their independently audited financial accounts with Companies House. This public agency makes the filings available to the public. From this source, we are able to gather financial statistics for almost all the clubs. Accounts in one form or another were filed for 95% of the clubs in our data, including data on wage expenditure for 85% of clubs. The table refers to clubs in the final sample, i.e. clubs for which we have at least 'some' years of financial information.

Table 3: Ability estimates at individual manager level by quartile in estimated ability

	Full sample			Ability quart 1		Ability quart 2		Ability quart 3		Ability quart 4	
Ability	**Obs.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Mean value	626	0.234	0.835	-0.792	0.800	0.123	0.121	0.495	0.113	1.116	0.455
Number monthly est.	626	36.3	43.3	15.7	22.0	35.7	37.8	52.8	52.3	40.9	46.4
Total obs. in data	626	141.4	175.5	64.3	102.6	141.8	152.9	206.6	212.5	152.8	185.2

Notes: Table depicts the mean ability estimate, number of monthly estimates and total number of observations at the level of the individual manager. Numbers refer to the average ability estimate for a manager across all observed periods. We provide statistics for the full sample and a breakdown by quartiles in estimated manager ability distribution.

Table 4: Count of 'substandard' experienced managers hired by division and comparison group

Division	# Exp. hires	# Sub-standard	% Sub-standard	Comparison group	Av. # entrants in comparison
All	755	192	25.4%	Mean entrant previous 5 years all divisions	96
1	142	24	16.9%		96
2	235	47	20.0%		96
3	181	56	30.9%		95
4	197	65	33.0%		95
All	755	247	32.7%	Mean entrant previous 5 years own division	25
1	142	53	37.3%		13
2	235	76	32.3%		28
3	181	48	26.5%		29
4	197	70	35.5%		27

Notes: Table shows the absolute and relative number of experienced managers who are substandard at the time of hiring. We use the entrants in the past 5 years and 10 years as the comparison group. Alternative ability measures and assumptions on the entrant comparison group lead to comparable results. See appendix A for more detail.

Table 5: Transition probabilities from 'substandard' to 'above standard' over duration of employment spell

		End of Spell		Observations	Comparison group
		Substandard	Above Standard		
Start of Spell	Substandard	72%	17%	192	Mean all divisions
	Above Standard	6%	94%	563	
	Substandard	67%	33%	247	Mean own division
	Above Standard	11%	89%	508	

Notes: Table displays the transition probabilities for a manager to move from substandard to above standard and vice versa over the course of an employment spell. Substandard managers are defined in comparison to the mean entrant in the past 5 years, either in all divisions or in the hiring club's division. At both the start and end of the employment spell, the ability estimate of the experienced manager is compared to the contemporaneous entrant distribution.

Table 6: LPM results for probability that experienced manager is hired as 'substandard'

Comparison	Substandard experienced hires vs. mean entrant in past 5 years				
Learning:					
Log games	-0.117***		-0.138***	-0.150***	-0.135***
	(0.013)		(0.015)	(0.015)	(0.015)
Log non-Eng. games	0.035***		0.043***	0.040***	0.043***
	(0.010)		(0.011)	(0.011)	(0.011)
Exp. < 20 games		Ref.			
Exp. 20-40 games		-0.109			
		(0.095)			
Exp. 40-80 games		-0.181**			
		(0.082)			
Exp. 80+ games		-0.397***			
		(0.067)			
Division dummies:					
1st division			Ref.	Ref.	Ref.
2nd division			0.031	0.034	0.037
			(0.044)	(0.044)	(0.044)
3rd division			0.136***	0.134***	0.149***
			(0.047)	(0.047)	(0.048)
4th division			0.150***	0.148***	0.152***
			(0.046)	(0.046)	(0.047)
Other controls:					
Log age			0.381***	0.472***	0.378***
			(0.127)	(0.134)	(0.132)
Foreign			-0.249***	-0.260***	-0.254***
			(0.091)	(0.092)	(0.093)
Log month number			-0.256***	-0.254***	-0.251***
			(0.047)	(0.047)	(0.048)
Constant	0.848***	0.595***	-1.954*	-2.800**	-2.159*
	(0.070)	(0.065)	(1.156)	(1.219)	(1.212)
Entry mode	No	No	No	Yes	No
Playing career	No	No	No	No	Yes
Month FE	No	No	Yes	Yes	Yes
Obs.	755	755	755	755	755
R-squared	0.097	0.074	0.152	0.163	0.161

Notes: Table reports regression results for a linear probability model where the dependent is an indicator equaling 1 if a rehire is substandard, 0 otherwise. The comparison group are the mean entrants in all divisions over the past 5 years. Regressions using other ability measures or comparison groups yield equivalent results and are available on request. The top panel includes measures of manager experience, the middle panel reports division dummies, the lower panel highlights other personal characteristics. The categories 'entry mode' and 'playing career' refer to the variables given in Table 1 under these headings. We do not report point estimates for these variables to aid readability. Bootstrapped standard errors are given in parentheses, significance is denoted by *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Summary ability and career length for counterfactual workforce

	# Manager-months	# individual managers	Mean career length (months)	Mean ability estimate	Bootstrap std. error
Replaced spells					
Actual	4,324	100	54.8	0.144	0.007
Counterfactual	4,324	271	16.4	0.511	0.013
Difference	4,324			0.367	0.014
Full sample					
Actual	22,697	596	40.9	0.428	0.004
Counterfactual	22,697	867	29.3	0.498	0.005
Difference	22,697			0.070	0.003

Notes: This table shows the results of a counterfactual, where we drop all employment spells in managerial careers, which occur after a manager has been hired as a substandard manager (even if the particular spell did not start with a substandard hire). We replace these spells by draws from the observed entrant ability distribution, again disregarding spells following a substandard hire. We show average career length and estimated worker ability for the subsample we replace (top panel) and the full sample (bottom panel). The level of observation is the individual manager-month, with the number of observations indicated in the first column. We perform 1000 bootstrap replications of this counterfactual to assess the significance of the difference between the actual and counterfactual samples.

Table 8: Linear regression results for probability of employment spell termination in month of analysis

	Spell Termination		
Ability estimate	-0.019***	-0.017***	-0.021***
	(0.003)	(0.003)	(0.004)
Dif. ability estimate over month		-0.025***	-0.028***
		(0.009)	(0.005)
Log tenure			0.004***
			(0.001)
Log age			0.067***
			(0.017)
Log experience			-0.007***
			(0.002)
Constant	0.061***	0.058***	-0.576***
	(0.002)	(0.002)	(0.156)
Observations	19,591	18,976	18,945
Month FE	No	No	Yes
Division FE	No	No	Yes
Year FE	No	No	Yes
R-squared	0.003	0.002	0.032
Adj. R2	0.003	0.002	0.029

Notes: Table reports linear probability model estimates where the dependent variable equals 1 if manager leaves the club in the month of analysis and 0 otherwise. Separate analyses using alternative ability measures and/or Cox hazard models yielded equivalent results and are available on request to the authors. We report bootstrapped standard errors in parentheses, indicating significance as follows, *** p<0.01, ** p<0.05, * p<0.1.

Table 9: Regression results for career progression after current employment spell

Dependent Variable	Ever rehired indicator			Career progress		
Ability est. at spell end	0.158***	0.090***	0.086***	0.335***	0.191***	0.186***
	(0.019)	(0.025)	(0.021)	(0.060)	(0.068)	(0.066)
Average win%		-0.008	-0.009		-0.027	-0.028
		(0.015)	(0.016)		(0.039)	(0.030)
Average goals pro		0.042***	0.044***		0.112***	0.118***
		(0.014)	(0.017)		(0.041)	(0.038)
Log games		0.059***	0.061***		0.144***	0.141***
		(0.013)	(0.012)		(0.026)	(0.034)
Log age		-0.587***	-0.596***		-1.456***	-1.608***
		(0.128)	(0.098)		(0.263)	(0.285)
Foreigner		0.223***	0.117*		0.390***	0.167
		(0.052)	(0.060)		(0.142)	(0.152)
Log month number at spell end		0.141***	0.126***		0.341***	0.305***
		(0.038)	(0.037)		(0.086)	(0.091)
Entry mode	No	Yes	No	No	Yes	No
Playing career	No	No	Yes	No	No	Yes
Month FE	No	Yes	Yes	No	Yes	Yes
Observations	1,250	1,244	1,244	1,250	1,244	1,244
(Pseudo-)R2	0.058	0.127	0.133	0.016	0.043	0.047
Estimation method	Linear Probability			Ordered Probit		

Notes: The left hand panel displays regression results for a linear probability model where the dependent variable is an indicator which equals 1 if a manager is rehired as a manager anywhere in England or abroad after the current employment spell ends, and 0 otherwise. The right hand panel shows an ordered probit model, where the dependent variable categories are defined as follows, 0 if the manager is never rehired; 1 if the manager is rehired in a lower division as where current spell started; 2 if the manager is rehired in same division as where current spell started, and 3 if the manager is rehired in a higher division as where current spell started. Foreign leagues are deemed equivalent to the English second division, apart from the big 4, who are deemed equivalent to division 1. The categories 'entry mode' and 'playing career' refer to the variables given in Table 1 under these headings. Point estimates are not reported to aid readability, but are available on request to the authors. We report bootstrapped standard errors in parentheses, with significance levels given as *** p<0.01, ** p<0.05, * p<0.1.

Appendix A: Results alternative ability measures and comparison groups

In this appendix, we describe a set of alternative manager ability measures, which we calculate in addition to the baseline measure reported throughout the main body of the paper.

First, we address potential feedback effects from game results to input values in the estimation of equation (5), which would primarily stem from performance-based bonuses in the players' compensation packages. To this end, we follow Peeters and Szymanski (2014) and rerun each model using past payroll and assets to instrument for current values (the resulting estimates are dubbed 'Base + IV' below). While this procedure corrects for contemporary feedback, we recognize it cannot fully alleviate concerns over endogenous input choices. Second, we estimate the rolling manager effects using additional controls for manager experience, i.e. the level, square and cube of log career games coached ('Base + experience' below). Third, we allow for the possibility that observations that are more recent provide more information on current managerial ability. We therefore re-estimate the model weighing observations by the number of days between the date of the game and the estimation month (see 'weighted base' below). We finally report results from a two-stage estimation procedure introduced by Jackson (2013) to allow for firm-worker match quality (see Jackson, 2013, Stanton et al., 2015 and Peeters et al., 2015 for more). In this approach, referred to as 'Spell' below, we estimate firm-manager spell effects in the first stage, and then split these into manager and firm effects using weighted least squares with inverse first stage standard errors as weights. This method forces a mean zero assumption on the match quality of all spells a manager has over his career in the data, but allows for differing match qualities among his spells. As an alternative to the second stage, we also report the manager's average spell effects ('Av. spell'), weighting the spells by their number of observations.

As a further robustness check we also introduce a cruder measure of managerial ability, 'add win%', which requires no formal econometric analysis. Here we define managerial ability, wp_{mp}, as the average winning percentage a manager m realizes at all teams i in division d relative to the historic winning percentage of those teams i at level d before manager m's arrival ($\overline{wp}_{-mid}$). More formally, we specify that

$$wp_{mt} = \sum_{g=1}^{g=n_p} \frac{1}{n_p}\left(w_{gmid} - \overline{wp}_{-mid}\right), \qquad (6)$$

where n_p refers to the number of games the manager has been active up until month p and w_{gmid} is the result of game g, coached by manager m at team i in division d expressed as 1 for a win, 0.5 for a draw and 0 for a loss.

In Table 10 we report the correlation between our two main alternative measures of manager ability (the baseline from the main analysis and Add win%) and the alternative specifications outlined above. The variants of the base methodology (Spell, Av. spell, IV, Experience, Weighted and combinations thereof) are all highly correlated with our baseline measure. While the cruder Add win% model has a lower correlation it is still in the range 0.6 to 0.7, which is striking given the variety of factors not controlled for in this naïve model. Moreover, both measure correlate highly (>0.5) with very crude performance measures, such as average win percentage and goal difference.

In table 11, we use alternative measures and comparison groups to replicate the main results in table 5 . In the top two panels, we vary the entrant comparison group, by replacing the average for the median and increasing its range from 5 to 10 years. In the bottom five panels, we report results for alternative estimates of manager ability, i.e. the add win%, the two-stage spell estimator, the baseline estimator with weights for recent performances and the base estimator including a polynomial of manager experience. All results point in the same direction, i.e. between 20% and 35% of all hiring events involves a substandard experienced manager.

Table 10: Correlation among ability estimates at spell level

Correlation Coefficient	Add win%	Spell	Av. Spell	Weighted Base	Base + exp.	Base + IV	Spell + IV	Weighted base + IV	Base + exp. + IV	Av. win %	Av. Goals pro
Base ability	0.678	0.900	0.824	0.978	0.986	0.966	0.841	0.950	0.950	0.606	0.502
Add win%		0.613	0.693	0.662	0.676	0.682	0.600	0.669	0.675	0.846	0.646

Notes: Table reports the correlation coefficient between different estimation methods for managerial ability and the baseline measures used in further analysis. We take the ability at the end of each spell as the unit of observation. Method names include 'IV' when wages and assets are instrumented with past values, 'exp.' when estimation includes experience polynomial, and 'Spell' when ability is estimated with manager-club spell dummies, which are decomposed in a 2^{nd} stage. We do not separately report significance, as this is always found to be below $p<0.01$.

Table 11: Presence of mediocre rehires for alternative ability measures and entrant comparison

Ability Estimate	Division	# Exp. hires	# Sub-standard	% Sub-standard	Entry comparison group	Av. # Entrants in comp.
Baseline	All	755	171	22.6%	Mean entrant previous 10 years all divisions	201
	1	142	22	15.5%		201
	2	235	40	17.0%		201
	3	181	50	27.6%		200
	4	197	59	29.9%		201
	All	754	216	28.6%	Median entrant previous 5 years all divisions	96
	1	142	27	19.0%		96
	2	235	53	22.6%		96
	3	181	63	34.8%		95
	4	196	73	37.2%		95
Add win%	All	819	164	20.0%	Mean entrant previous 5 years all divisions	101
	1	150	27	18.0%		102
	2	250	41	16.4%		102
	3	202	43	21.3%		101
	4	217	53	24.4%		101
	All	819	205	25.0%	Mean entrant previous 5 years own division	27
	1	150	23	15.3%		13
	2	250	69	27.6%		29
	3	202	41	20.3%		30
	4	217	72	33.2%		30
Spell	All	755	184	24.4%	Mean entrant previous 5 years all divisions	96
	1	142	19	13.4%		96
	2	235	46	19.6%		96
	3	181	55	30.4%		95
	4	197	64	32.5%		95
Weighted baseline	All	755	192	25.4%	Mean entrant previous 5 years all divisions	96
	1	142	23	16.2%		96
	2	235	48	20.4%		96
	3	181	55	30.4%		95
	4	197	66	33.5%		95
Experience baseline	All	755	178	23.6%	Mean entrant previous 5 years all divisions	96
	1	142	22	15.5%		96
	2	235	49	20.9%		96
	3	181	52	28.7%		95
	4	197	55	27.9%		95

Notes: See table 5 notes.

Appendix B: Variable definitions

Name	Definition	Type
Table 1:		
• Games observed	Number of games manager is observed in the data, not taking into account potential missing variables.	Continuous
• Months observed	Distinct calendar months in which manager is observed in data, not taking into account potential missing variables.	Continuous
• Av. Age	Average age of manager in years over all his observations in data.	Continuous
• Av. Exp.	Average experience in games (both in England and abroad) of manager over all his observations in data.	Continuous
• Av. Eng. Exp.	Average experience in games (confined to England) of manager over all his observations in data.	Continuous
• Foreigner	Equals 1 if manager is not UK or Irish national, 0 otherwise.	Indicator
• Player-manager	Equals 1 if manager started manager career as player-manager, 0 otherwise.	Indicator
• Other man. exp.	Equals 1 if manager had other management function in football before 1st employment as manager, 0 otherwise.	Indicator
• Intern hire	Equals 1 if manager had other management function in same football club as 1st employment as manager, 0 otherwise.	Indicator
• Division	Division in which manager obtained first employment, ranked 4 (lowest) to 1 (highest).	Categorical
• Play prof.	Equals 1 if manager played as professional player, 0 otherwise.	Indicator
• Play big 4	Equals 1 if manager played in any of big 4 leagues (1st div. in Eng, Ger, Spa, Ita), 0 otherwise	Indicator
• Num. Eng. Team	Counts number of English clubs manager has played for.	Continuous
• Ex-player club	Equals 1 if manager is employed by club that he also played for, 0 otherwise.	Indicator
• International	Equals 1 if manager played for his respective national team, 0 otherwise.	Indicator
Table 2:		
• Years observed	Number of seasons club appears in data.	Continuous
• Wages	Total wage cost (incl. tax and social insurance) of club over season, from financial accounts.	Continuous
• Revenue	Total revenues of club over season, from financial accounts.	Continuous
• Fixed assets	Total book value of fixed tangible assets of club over season, from financial accounts.	Continuous
• Pos. net assets	Equals 1 if clubs has positive net assets over season, from financial accounts, 0 otherwise.	Indicator
• Pre-tax profits	Profits before taxation of club over season, from financial accounts.	Continuous
• # insolvencies /club	Number of insolvency proceedings club went through over sample period.	Continuous
• Goals pro	Total goals scored by club over all games in season.	Continuous
• Goals against	Total goals conceded by club over all games in season.	Continuous
• Goal Difference	Difference between goals scored and conceded by club in season.	Continuous
• Points	Number of ranking points obtained by club over season.	Continuous
• Win %	Percentage of games won by club over season, draw counted as half a win.	Continuous

• Number hires	Number of hires club has made over sample period.	Continuous
• Av. hires/year obs.	Number of hires club has made over sample period divided by number of years club is observed.	Continuous
• Av. manager tenure	Average tenure of manager at club over sample period.	Continuous
• # entrant hires	Number of times club has hired novice manager over sample period.	Continuous
• # foreign hires	Number of times club has hired foreign manager over sample period.	Continuous
Table 6:		
• Log games	Logarithm of total number of games manager has managed prior to hiring.	Continuous
• Log non-Eng. games	Logarithm of total number of games manager has managed in England prior to hiring.	Continuous
• Exp. < 20 games	Equals 1 if manager has less than 20 games (i.e. half season) experience prior to hiring, 0 otherwise.	Indicator
• Exp. 20-40 games	Equals 1 if manager has between 20 and 40 games experience prior to hiring, 0 otherwise.	Indicator
• Exp. 40-80 games	Equals 1 if manager hasbetween 40 and 80 games experience prior to hiring, 0 otherwise.	Indicator
• Exp. 80+ games	Equals 1 if manager has more than 80 games experience prior to hiring, 0 otherwise.	Indicator
• Log age	Logarithm of manager age, expressed in days, at day of hiring.	Continuous
• Foreign	Equals 1 if manager is not UK or Irish national, 0 otherwise.	Indicator
• Log month number	Logarithm for a counter of the month of hiring within the data, August 1974=1, May 2011=380.	Continuous
Table 8:		
• Ability estimate	Ability estimate of manager in month of analysis.	Continuous
• Dif. ability estimate over month	Difference between ability estimate in month of analysis and ability estimate in previous month.	Continuous
• Log tenure	Logarithm of the number of games in current employment spell, i.e. manager-club pairing.	Continuous
• Log age	Logarithm of average manager age, expressed in days, in month of analysis.	Continuous
• Log experience	Logarithm of total manager experience, expressed in games, in month of analysis.	Continuous
Table 9:		
• Ability est. at spell end	Ability estimate of manager in month spell ends.	Continuous
• Average win%	Career average win percentage obtained by manager in month spell ends.	Continuous
• Average goals pro	Career average goals scored obtained by manager in month spell ends.	Continuous
• Log games	Logarithm of number of games manager has managed up until month spell ends.	Continuous
• Log age	Logarithm of manager age in days in month spell ends.	Continuous
• Foreigner	Equals 1 if manager is not UK or Irish national, 0 otherwise.	Indicator
• Log month number at spell end	Logarithm for a counter of the month spell ends within the data, August 1974=1, May 2011=380.	Continuous