TI 2017-116/VII
Tinbergen Institute Discussion Paper

The Survival of Mediocre Superstars in the Labor Market

Revision October 2021

Thomas Peeters[1]
Stefan Szymanski[2]
Mark Terviö[3]

[1] Erasmus University Rotterdam, Tinbergen Institute, ERIM, ECASE
[2] University of Michigan
[3] Aalto University

The Survival of Mediocre Superstars in the Labor Market

Thomas Peeters[1]
Erasmus School of Economics, Tinbergen Institute, ERIM and ECASE
peeters@ese.eur.nl

Stefan Szymanski[2]
University of Michigan, Department of Sport Management, School of Kinesiology

Marko Terviö[3]
Aalto University, Department of Economics

Abstract

We argue that liquidity constrained firms face strong incentives to hire experienced, but low ability workers instead of novice workers with higher upside potential. Using four decades of high-frequency information on worker performance in a 'superstar' labor market allows us to estimate the revealed ability of experienced workers at the time they are hired by a new firm. More than one fifth of these hires are "substandard" in that the revealed ability of the hired experienced worker lies below the mean ability of recent novices. Even more hires (around 40 percent) are "mediocre", as their ability falls short of the hiring threshold that maximizes the long-run average ability of the active workforce. Replacing mediocre hires by novice workers would increase the average ability of the workforce by 0.1 standard deviations.

JEL Codes: M51, J63, J24, Z22

Keywords: hiring, labor market entry, superstars, football

1. Introduction

The option value of employment relationships is critical for efficient turnover decisions in labor markets (Jovanovic, 1979). Hiring a run-of-the-mill worker, whose ability is well known, is socially inefficient if it is possible to hire someone else, who is equally able in expectation, but could turn out better or worse. In this case, upside risk is more valuable than downside risk, because successful workers have longer careers than unsuccessful ones. Hence, it is socially efficient to hire riskier job candidates. As argued by Terviö (2009), a problem may arise when worker performance is readily observable to competing employers, e.g. in creative industries such as music and sports. Here, the finders do not get to keep the stars they discover - at least not at a wage that would leave enough rents to compensate employers for their initial investment. When differences in ability entail large differences in economic value, it is also unrealistic to expect entry-level workers to compensate firms for the full value of the chance to be discovered. In this case, employers rationally adopt a lenient threshold for retaining experienced workers. This allows “mediocrities”, workers whose ability is already revealed to be below the threshold for efficient retention, to survive in the labor market. As a result, average worker ability in the industry is inefficiently low (Terviö, 2009).

In this paper we extend the model of Terviö (2009) to incorporate the possibility that the firm faces a liquidity constraint and risks becoming insolvent (bankrupt) when she hires a worker of too low ability. This is a well attested problem in the industry we study (Szymanski, 2017), but also in the wider economy. We show that insolvency problems in this case further reduce the incentive to hire novices, up to the point where firms may hire experienced workers whose ability lies below the mean ability of recent novices, even absent concerns about poaching. Firms will prefer an experienced worker whose known ability guarantees the survival of the firm at the expense of a potentially more lucrative novice, who carries a small risk of bankrupting the firm.

We explore the implications of our reasoning in a dataset containing 38 years of high frequency data on worker performance drawn from one market for “superstar” workers, the labor market of English football managers. We use the data to (re-)estimate each active worker’s ability after each calendar month, based solely on contemporaneously available performance information. For each instance where we observe an employer hiring worker with prior experience, we compare the estimated ability of the selected worker to the ability of recently hired ‘novice’ workers without prior experience. We define two thresholds for this comparison. We label around one fifth of the hires in our analysis “substandard”, meaning that the revealed ability of the hired experienced worker is lower than the mean ability of recent novices. As in Terviö (2009), we use the term “mediocre” to denote rehired workers whose revealed ability lies below the threshold that would maximize the average ability of the active workforce in the labor market. More than forty percent of the hires we analyze qualify as mediocrities. These findings clearly imply that employers retain too many low ability experienced workers and we show that average worker ability indeed falls short of reaching its full potential.

The information structure of the labor market we study is such that information on worker ability is extremely limited before workers enter their first employment spell. Still, as workers gain experience, their ability is quickly revealed to the market as a whole. There is a publicly observable and commonly acknowledged benchmark of success in the industry and the contribution of the worker can be separated from the contribution of other inputs thanks to publicly available, audited financial statements. In combination with the inability of workers to either finance their own entry into the labor market or credibly commit to long-term contracts with their initial employer, these circumstances closely resemble the setting of Terviö (2009). At the same time, periods of financial distress are commonplace among the employers in this industry (Szymanski (2017)). Using a simple theoretical hiring model, we show that this feature exacerbates their tendency to prefer experienced workers of known low ability over novices with unknown, but potentially higher ability.

We empirically study the efficiency of hiring decisions by comparing the ability of rehired experienced workers with a counterfactual of hiring novices in their place. To control for input use in the estimation of worker ability, we use individual effects estimators à la Abowd, Kramarz, and Margolis (1999). While the market we study represents a classic example of team production, the football manager's distinct role allows us to identify the contribution of this specific individual. We adopt an updating estimation procedure to capture changes in perceived worker ability in reaction to observed success and failure. We estimate the model using only observations up to the time of hiring, so our ability measure is consistent with what potential employers could have observed at each point during a worker's career. Given the performance of the experienced worker to date, we can compare his estimated ability to the expected ability of a novice, which we assume to be drawn from the estimated abilities of actual novices in the previous ten years. When we analyze the incidence of hiring mediocre workers, it appears that firms with less favorable finances and lower productivity are more likely to hire a mediocre experienced worker instead of a novice. Mediocre experienced hires are relatively older and less experienced workers.

An important contribution of our paper is that we analyze hiring efficiency in a job, which is both very high profile and deemed to be of great significance to consumers who buy the product. Previous studies of inefficient hiring have used data from online labor markets where there is much less at stake. 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 consider the value this information has for later hiring decisions. While our data covers entire careers, some of which last several decades, we do not have a source of exogenous variation in hiring policies, so the nature of our empirical findings is necessarily more descriptive. Our focus on liquidity constraints echoes the recent work of Caggese et al (2019) who show that financing constraints lead firms to fire too many short tenure workers with large upside potential at the benefit of longer tenure workers with less scope for future productivity gains. However, the intuition in their study starts from firing costs, whereas our focus is on hiring instead.

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 a worker's career progresses. The information revelation mechanism in the industry we study is similarly symmetric. Without experience, neither the firm nor the would-be workers can know how well they are likely to perform. Ex post, the public nature of performance means the ability estimates of experienced workers are public information. Hence, we abstract from the possibility that firms can strategically conceal or reveal information on their workers' ability, as analyzed by Strobl and Van Wessep (2013). Most of the subsequent empirical literature has focused on the turnover of employees between firms in an industry investigating how firms learn about each other's employees. For example, organizations may leverage the social ties of their existing workforce to learn about job market candidates (Sterling, 2014) or use hiring intermediaries (Stanton and Thomas, 2016). Our focus is slightly different: we look at the selection of workers into 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 into and out of the industry.

As advocated by Kahn (2000) and Palacios-Huerta (2016), in our empirical analysis we use rich and long running individual performance data from the professional sports industry to draw lessons that are relevant to other labor markets, where data on worker performance is less readily available. In doing so, we contribute to the emerging strand of literature that takes sports data to look at questions in organizational and labor economics.[4] Within this literature, the performance and subsequent dismissal of football managers has received quite some attention (see e.g. Bridgewater, Kahn and Goodall (2011), Hope (2003), Van Tuijl and Van Ours (2016)), but, to the best of our knowledge, we are the first to investigate the selection of novice workers into this profession.

In the remainder of this paper, we first set out a theory model, which introduces firm liquidity constraints into the framework of Terviö (2009). Then we briefly sketch the institutional environment of the labor market for professional football managers in England. In the next section, we explain the dataset and empirical methodology we employ, followed by our main empirical results. We review the main contributions of our work in the final section.

2. A model of hiring with firm liquidity constraints

In this section, we present a formal model that shows why we can expect liquidity constrained employers to be more eager to hire experienced workers, in a way that lowers the ability of active workers in the industry.

We use the model to examine how profit-maximizing firms hire key workers, when they face a choice between inexperienced and experienced workers whose expected abilities may differ. The ability of inexperienced workers 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 worker ability can only be inferred after observed job performance at a firm. The number of real job opportunities is scarce, and there will always be more potential good workers than could ever be employed. So, there is always a pool of available inexperienced workers to recruit from. To understand the key economic trade-off here, notice that an employment spell produces not just current output, but also information about a worker's ability, which is useful for future output. It is therefore socially inefficient to hire an experienced worker who is known to be only slightly better than a novice 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, 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 workers 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 workers to commit to long-term wage contracts. If successful workers are easily poached or have their wage quickly bid up by competing firms, then a firm that discovers a new star worker 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 workers (a) cannot commit to long-term wage contracts and (b) are liquidity constrained, because they could otherwise "buy the firm".

The second problem arises when firms face a liquidity constraint as well. This means that a bad "draw" from the worker 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 worker 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 worker has no positive value for the current owners. A liquidity constrained firm will therefore give more weight to the downside risk in worker ability and may hire a worker of known low quality over a riskier novice, even if the latter is expected to perform better. This is the key intuition in our model.

2.1. Model setup

Consider a firm with revenue equal to the ability of the worker. There is a population of untried workers with ability drawn from a known continuous distribution. Worker ability θ becomes known after one period of work, and 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 workers 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 workers, so novices are held to their outside wage, which we normalize at zero.

The firm only faces a real decision in periods when it has an incumbent employee: 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 workers when their ability is above ψ and otherwise hire a novice.

2.2. 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. This is necessarily the case when it employed an experienced worker in the previous period. If it employed a novice in the previous period, then it knows the ability θ of its incumbent 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 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 of ability θ can therefore be written as,

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

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 π(θ) is from (1):

$$V_0 = \Pr(\theta \ge c)\,(\mathrm{E}[\theta|\theta \ge c] - c) \\ + \delta\{\Pr(\psi > \theta \ge c)\,V_0 + \Pr(\theta \ge \psi)\,(\mathrm{E}[\theta|\theta \ge \psi] - c + \delta V_0)\} \tag{2}$$

The last term takes into account that if the worker turns out to have θ > ψ, 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] and that c<0.5 to focus on the interesting case. Then expression (2) becomes

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

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

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

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

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 workers would be deemed "mediocre" and not be rehired, and certainly no "substandard" workers, whose ability is below the mean entrant, would be rehired.

2.3. Discussion

In Figure 1, we analyze how the firms' optimal hiring policy (4) depends on the severity of the liquidity constraint c and on the discount factor δ. We plot the optimal rehiring threshold ψ_c as a function of c for selected values of δ. For example, using $c = 0.1$ and $\delta = 0.95$ yields $\psi_c \approx 0.32$, so the firm would hold on to incumbents who are substantially worse than population average - even though novices, 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 workers in order to avoid the risk of disastrously bad novices.

< Insert Figure 1 around here >

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 $\delta \rightarrow 1$. 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 $c > 0$, a different type of "investment" incentive appears. The hiring of a known mediocre or substandard worker 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. For firms with a modest liquidity constraint and discount factor, the privately optimal hiring policy is too lenient relative to social optimum, and below the population mean for a wide range of parameters. With very low δ and high c the private threshold can also be higher than is socially optimal, but this implies extremely high discounting, which is not likely to be realistic.

The most natural case is one where liquidity constrained firms find it in their interest to hire and retain experienced workers that are in expectation worse than novices. This turns out to be empirically relevant in our analysis of the English football industry, where financial constraints are commonplace (Szymanski, 2017), and, as we argue below, about one fifth of all hired experienced workers has an estimated ability below the mean ability of recent novices. Note that this is a finding that could not be explained within the existing framework of Terviö (2009).

3. The labor market for football managers in England

Professional football is the world's most popular spectator sport. The English professional leagues are amongst the best known in the world and the performance of each club is publicly scrutinized in immense detail. The football manager is hired by the club to oversee the on-field performance of the team. In England the manager is also typically responsible for hiring and firing players, training, and motivating the players under contract, and setting the team strategy for each game. There can be few, if any, professions, which are the subject of such intense interest.

The principal measure of a manager's success is how many games he wins. This characterization is supported by commentary in the media, as well as the testimony of employers and the managers themselves. It is also well supported in the academic literature (e.g., Van Tuijl and Van Ours (2016)). Of course, wins are observable. Still, even if winning is a natural measure of success for a football manager, simply counting the number of wins does not do justice to the ability of the manager. First, the margin of victory can be large or small and hence the difference in goals scored is usually a more precise measure of the outcome. Second, there are two factors that play a significant role in determining the probability of winning or scoring – whether the team plays at home or away and the relative resources employed by both teams. In football the difference in resources, measured either by club revenues or wage bills, is enormous. Among the one hundred or so English football clubs active over the last forty years, the difference between the largest spending and lowest spending team in any one year has ranged between a ratio of 10 to 1 (early in the sample) and 200 to 1 (in recent years). Thus, we estimate worker ability based on wins, but we adjust for resources and home advantage.

Our dataset contains information on approximately 75,000 games played in English professional football over the period 1974-2011. A typical league season comprises about forty games for each team. 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 significantly changed. The rules of the game itself have also changed very little. Winning games has always been the yardstick of performance.

Each football club has a board of directors and CEO responsible for the day-to-day management of the business side of the club. It is very rare for a football manager to move beyond the sporting aspects of the club's activities and become a member of the board of directors. Hence, in business terms their job profiles are similar to those of middle managers. Yet, the public prominence of football means that many of them become household names.

Most managers are recruited from the ranks of former players, and a few even become managers while still playing. In most cases, 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. Occasionally, high profile former players get the chance to progress more rapidly through the system. Formal qualifications were not required before 2003. Since then managers in the top division have to obtain a UEFA Pro License, which takes about a year of part-time study to acquire.[5] Exemptions from this requirement are still 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. In-season manager dismissals are common. 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 our game level data. For example, the 37% of managers with a career length under 2 years represent only 7% of all game observations.

Unlike in the US major leagues, there are no systematic public records of the salaries and contractual arrangements for individual managers in English football. However, media reports suggest that almost all employment contracts are short-term, spanning one or two seasons. Even when a club employs the same manager for an extended period of time, the board typically rolls over a series of short-term contracts, rather than offering a single long-term deal.[6] An important motivation for clubs to avoid long-term contracts is that this would expose them to significant firing costs, which is a real concern because employment spells frequently end with the club firing the manager before the contract ends. The labor market for managers is not covered by the “transfer system”, which applies to football players. In contrast to the practice for players, clubs can poach a manager from another team without having to pay a high compensation to their current employer. As a result, clubs cannot force their manager to stay when a better employment offer comes along, and, even if they could, it would be unwise to hold on to a manager set on leaving, because important managerial tasks cannot be contractually enforced. This situation makes it unlikely that clubs can overcome the inefficiencies described in Terviö (2009) through contractual solutions. After all, this would require contracts that credibly commit novice workers to their initial employer for the majority of their (potential) career (Terviö, 2009; p. 838).

There are 92 English professional clubs competing in each season, which play in four hierarchical divisions of roughly equal numbers (currently 20 in the top division, and 24 in each of the 3 lower divisions).[7] Each club plays every other club in its division twice (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). This system of promotion and relegation is present in football leagues around the world and differentiates its organizational structure from the professional sports leagues in the US. 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 in the American major leagues, mechanisms to promote revenue sharing and wage controls are very limited.

Over the last 40 years the primary components of revenues and costs in football have not changed much. Revenues come 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.[8] Because broadcast revenues are concentrated in the top division this trend has further increased inter-divisional inequalities. Producing commercial football games requires two essential assets: a stadium (which, in England, clubs typically finance and own themselves) 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 (e.g., Szymanski and Smith, 1997; 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 division 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 four to 1 in the seventies to over 20 to 1 in recent years.

4. Data and empirical methods

4.1. Dataset

In most activities, performance measurement is difficult and often imprecise, but in football, team performance is a matter of official record. A manager’s career can stretch over hundreds and possibly thousands of games, and after each game is played there is no uncertainty about its outcome measured by goals scored and conceded, which in turn determines the result (loss, tie 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.

< Insert Table 1 around here >

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 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, and (e) if they represented their national team.

We provide more detailed variable definitions Appendix A. We present the sample mean and standard deviation for these variables, and the mean within quartiles based on total career length. There is a very large variance in the experience levels of football managers. The mean experience at the end of a worker's 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 in our sample 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 they played for between three and four clubs in England in their careers. 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).

< Insert Table 2 around here >

All professional English football clubs take the legal form of limited liability companies. We retrieve financial information, such as revenues, wages and profits from the audited accounts of these companies filed in a public register at Companies House. Table 2 summarizes the data on the clubs, which employ the managers we observe. 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, 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 that 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 workers but are more likely to hire a foreign worker.

4.2. Assessing worker ability

We derive our primary estimate measure of worker ability from a regression model explaining the goal difference, y_{ijt}, in a game between clubs i and j in season t, based on a vector of inputs for each club (X_{it} and X_{jt}), a set of firm effects (γ_i and γ_j), and worker (manager) effects (μ_m and μ_k), for manager m of team i and manager k of team j respectively. Our baseline model is:

$$y_{ijt} = \beta_d X_{it} - \beta_d X_{jt} + \gamma_i - \gamma_j + \mu_m - \mu_k + \varepsilon_{ijt}. \qquad (5)$$

The subscript d indicates that the coefficients on variable inputs may differ by the division in which the clubs play. The first term in vector X_{it} is a dummy denoting whether team i plays at home, which is relevant because of the persistence of home advantage in professional team sports (see e.g. Garicano et al., 2005). As in Peeters and Szymanski (2014), we control for the player inputs by means of the logarithm of the total wage bill paid out by club i in season t. We further include the logarithm of 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 identify the worker and firm effects in equation (5) mimics the identification strategy of Abowd et al. (1999).[9] Using the routine of Cornelissen (2008), we first determine which clubs are connected by moving workers. We then drop all game observations pertaining to clubs outside the largest connected network. Note that both clubs i and j have to be part of the connected network for the game to be included. We finally estimate the worker and club effects for month p by entering individual dummies in the linear regression of equation (1). We use each game in the dataset twice, once from each team's perspective. This avoids the need to impose linear restrictions on the worker (club) dummies to ensure that these are estimated exactly opposite when the worker (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 appendix B, we consider two alternative game outcome measures, the win percentage (win%), defined as 0 for a loss, 0.5 for a tie and 1 for a win, and the increase/decrease in Elo rating (see Hvattum and Arntzen, 2010). We also employ a host of alternative approaches to estimate the baseline ability we specify in equation (1). These either (a) include controls for firm-worker match quality, (b) put more weight on recent, rather than historical performance, or (c) condition on a polynomial of working experience. We further examine whether the worker-firm network is sufficiently connected to avoid biases resulting from limited worker mobility (Andrews et al., 2008; Jochmans and Weidner, 2019). In the interest of clarity, we only present results using the baseline ability estimate from equation (1) for the majority of our analyses. We refer to appendix B and C for detailed information and results of these robustness checks and assumption tests.

In addition, we calculate a cruder measure of worker ability, 'add win%', 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 worker ability is the capacity to funnel resources to improvements in the playing squad. While simplistic, this approach avoids potential econometric concerns about our model specification and estimation method. Here, we define worker ability, wp_{mp}, as the average winning percentage a worker m realizes at all teams i in division d relative to the historic winning percentage of those teams i at level d before worker m's arrival ($\overline{wp}_{-mid}$). Formally, this means we calculate:

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

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

4.3. Estimation Algorithm

Unlike the stylized two-period model we sketched earlier, worker careers last many periods and ability is not instantly visible at a single point in time. Employers learn about ability over the worker's career. In our setting, this process is explicit, as each game publicly reveals information on the ability of the managers involved. In our empirical analysis, we compare the revealed ability of experienced workers at the time they are rehired by a club to the ability of contemporary novices, which the club could have hired in their place. This implies the need to estimate each rehired worker's ability, based solely on information revealed before the firm makes its hiring decision.

To achieve this, we divide the dataset in 380 calendar months, labeled p, running from August 1973 to May 2011.[10] We do not include June and July, as during these months clubs do not play official league games. We then run through the dataset from month 100 (i.e. May 1983) to 380 (May 2011), to estimate the ability of each incumbent worker at the end of each month. In other words, at the end of May 1983, we estimate the ability of each worker 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 worker ability taking the additional four or five games played in August into consideration. We repeat the process for September 1983, and so on until May 2011. We start the algorithm in May 1983, such that the first 10 seasons in the data may function as a learning period. We ensure that our estimates are comparable across periods by fixing the identity of the reference worker to which all other worker effects are scaled, i.e. the worker with a 'zero' effect is kept the same in each period.[11]

Through this estimation algorithm, we obtain a profile of worker ability estimates over the worker's career, rather than a single, end-of-career estimate. We denote this baseline estimate of worker ability by $\hat{\theta}_{mp}$ for worker m in month p. For our analysis of mediocre hiring, we use the ability estimate obtained at the end of the last calendar month before the hire to proxy the assessment of the hiring firm at the time of the hire. As worker ability is only observable during employment, we focus our attention on instances where we observe a firm hiring a worker, who has at least had one previous employment spell in the data. We dub these workers 'rehired' workers.

To proxy the ability of potential novices we generate an empirical equivalent for the novice ability distribution. For each month p we collect all contemporaneous ability estimates for workers who entered their first employment spell in the English managerial labor market during the hundred months (ten years) leading up to month p.[12] We take both active and inactive workers in period p into account to avoid selecting on labor market survival.[13]

4.4. Defining mediocre and substandard hires

In the model we set out above employers adopt a hiring policy, where they hire an experienced worker instead of a novice, if her revealed ability $\hat{\theta}_{mp}$ at the time of hiring exceeds a profit-maximizing threshold value $\tilde{\theta}_{\pi pi}$. Crucially, the threshold adopted by profit maximizing employers lies below the efficient threshold, $\tilde{\theta}_{\omega p}$, which maximizes the ability of active workers in the labor market. Whereas Terviö (2009) derives the socially optimal threshold analytically, here, we need to infer it based on the ability estimates we obtain. Evidently, the efficient threshold depends on the ability distribution of potential novices. A first threshold we therefore consider is the expected ability of novices in the comparison group, i.e. the mean ability of all novices in the ten years leading up to the time of hiring. We dub rehired experienced managers, who fail to reach this level as "substandard" hires and the corresponding threshold value as the "substandard threshold".

While this might seem like a natural choice, the asymmetry of the upside and downside potential of each novice hire implies that the truly "socially optimal" threshold may lie above this value. This also implies that the threshold should depend on the discount rate of the social planner, as this will determine the weight adhered to (future) upside potential. For simplicity, we adopt the point of view of a social planner with a discount rate of zero, so the optimum coincides with maximizing the steady state value. We refer to hires who fail to clear this bar as "mediocre" hires and to the corresponding threshold as the "mediocrity threshold". Note that the two thresholds we use represent two limiting cases, one where a social planner places no value on future realizations of the work force ability and one where the planner is infinitely patient.

To identify the "mediocrity" hiring threshold, we introduce a procedure, which calculates the average ability of active workers in the labor market for a range of different counterfactual hiring thresholds, $\tilde{\theta}_p$. Under the counterfactual policy, we first consider all contemporaneous estimates of worker ability $\hat{\theta}_{mp}$ for all workers *m* in months *p* when they were actually managing a club. We then identify the workers whose estimated ability $\hat{\theta}_{mp}$ falls below $\tilde{\theta}_p$, the threshold under consideration, at any instance where they got rehired. Next, we truncate the careers of these workers by deleting their employment spells after the rehire where their estimated ability $\hat{\theta}_{mp}$ was below $\tilde{\theta}_p$. In other words, the counterfactual assumes that workers failing to meet the threshold once, were never rehired. We then fill the resulting reduction in worker-months by adding counterfactual ability estimates, which we randomly draw from the observed distribution of estimated abilities in the remaining population of novices. Note that the composition of this group also varies depending on the analyzed threshold and includes the initial spells of workers who may be rejected once their ability is revealed. We finally calculate the average of the contemporaneous ability estimates in this counterfactual sample for threshold $\tilde{\theta}_p$.

Using this procedure, we evaluate the average ability of the counterfactually active workers for a range of hiring thresholds. We start at the substandard threshold, where $\tilde{\theta}_p$ equals the mean novice ability and then gradually increase the revealed ability required of rehired workers until we arrive at the threshold which maximizes the average ability in the counterfactual sample of active workers. This procedure allows us to find the mediocrity threshold $\tilde{\theta}_{\omega p}$. To mitigate the impact of random sampling, we repeat this procedure 200 times and use the average threshold over all replications as our estimate of the mediocrity threshold. In a final step we normalize worker ability by subtracting the average ability in the observed worker population and dividing the resulting term by its standard deviation. As such, our unit of measurement for an increase in worker ability is the standard deviation of ability in the original population of active workers.

It may be argued at this point that league competition is a zero-sum game and hence there is no gain from increasing the average ability of the active workers. In addition, we cannot infer to what extent firms are risk averse in hiring. It is therefore unclear how we should weight foregone future ability against increased certainty about current ability in the social optimum. We fully recognize this caveat, which is inherent in using on-field performance to construct a measure of worker ability. Still, increasing managerial ability in terms of on-field expertise is a stated target of both the leagues (the Premier League and the Football League) and the national governing body, which regulates them (the Football Association). Thus, from the point of view of these organizations one can claim that the hiring policies of the individual football clubs are inefficient. Of course, this is not the same as saying that high worker ability in the English football industry is socially optimal on a global scale.

5. Empirical Results

5.1. Worker ability

In Table 3 we show summary statistics for our ability estimates at the level of individual workers. 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 -0.302 with a standard deviation of 1.282. This may seem odd given our standardization procedure, which subtracts the population average from the raw estimates. However, the table reports career averages, which gives relatively more weight to workers who with accumulate fewer monthly observations and tend to have lower ability. This is consistent with employers using the information provided by the market to screen employees to retain those who are more able, and who therefore enjoy longer tenures.

< Insert Table 3 and Figure 2 around here >

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 long-standing workers. These are Arsène Wenger, one of the all-time great managers of top club Arsenal; Micky Adams, who managed a string of clubs at various levels (among others, Fulham, Brentford, Leicester City and Coventry City); and Steve McMahon, who managed lower tier clubs Swindon and Blackpool. 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 worker ability estimates for workers with the same experience, expressed in months since labor market entry. 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. During his tenure at Arsenal Football Club, Wenger was widely considered one of the most successful managers in England. Micky Adams' ability estimates are much closer to the median of the workforce, which may have helped him to return to a new 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 losing his first job (with Swindon Town).

5.2. How many hires of experienced workers are mediocre and substandard?

Now we turn to comparing the ability estimates of rehired experienced workers to the distribution of novice abilities. In the full sample, we have 755 spells of experienced workers for which we can calculate the baseline ability measure at the time of hiring. In Figure 3, the x-axis measures these estimated worker abilities scaled by the mean novice ability at the time of hiring. Only the hires plotted to the right of the y-axis are therefore expected to improve team performance relative to a novice. Yet, our estimates suggest that around one fifth of all hires fall below this modest threshold. Figure 3 also depicts our preferred 'mediocrity' ability threshold, which maximizes the average ability in the leagues over our sample period. It is clear that a significantly larger portion of the distribution fails to reach this more ambitious threshold.

< Insert Table 4 and Figure 3 around here >

Panel 1 and 3 in Table 4 detail the numerical results for our baseline ability measure, while panel 2 and 4 focus on the added win percentage, as described by equation (6). In the top two panels we compare rehires to the "mediocre" threshold. On this basis, we estimate that 326 (43%) of these rehires, involving 192 different individual workers were inefficient based on our mediocrity threshold. We also find that the proportion of mediocre rehires is higher in the lower tiers of competition – 53% and 46% in the lowest two tiers versus 33% in the top tier. This is consistent with the idea that lower division clubs have fewer opportunities to hire high ability experienced workers, such that they more often face the choice between a mediocre experienced worker and a novice. The pattern remains broadly the same for the added win% measure: 44% of rehires fall short of to the threshold and substandard workers are more likely to be hired in the lower tiers. For around twenty percent of all hires, the ability estimate of the rehired worker is below the mean ability of novices. Again, the pattern across tiers shows a higher probability of inefficient rehires in the lower tiers. The results using the added win% measures draw a similar picture. We further scrutinize the robustness of these results in Appendix B using alternative ability measures, estimation methods and comparison groups.

5.3. How much less talented is the active workforce because of mediocre hiring?

A central result in our model is that mediocre hiring may lower the talent level in the population of active workers. To measure the extent of this aggregate effect, we construct a distribution of worker ability following the counterfactual policy rule of never re-hiring workers, who are revealed as mediocre at any previous hire during their career. We follow a similar procedure to the one we developed to identify the mediocrity threshold. We consider the estimates $\hat{\theta}_{mp}$ for all workers m in months p when they were actually employed by a club but truncate the careers of mediocre workers to their employment spells before they were rehired as a mediocre experienced worker. As such, we assume counterfactually that mediocre workers were never re-hired.[14] We fill the resulting reduction in worker-months by adding counterfactual ability estimates, which we draw from the observed distribution of estimated abilities in the non-mediocre novice population. We then calculate the average ability in the actual and counterfactual workforce at the level of the individual worker-month, i.e. a worker'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. To aid interpretation, we normalize the average ability in the observed active workforce to zero and scale these abilities by the standard deviation in the original population of estimated abilities. We perform 1000 bootstrap replications of this procedure drawing a new sample in each iteration to assess the significance of the effects we uncover. [15]

< Insert Table 5 around here >

In the top panels of Table 5 we consider two counterfactuals. First, we exclude 7,096 worker-months because they belong to the careers of workers, who fall below the mediocrity threshold. The mean estimated ability of the counterfactual replacements is significantly higher than the original they replace: 0.179 compared to -0.175, a difference of 35% of a standard deviation in the original ability distribution. Second, we consider the effect of only replacing those experienced workers rehired with ability below the mean novice. This is a much smaller group (4,227) whose revealed average ability is much lower (-0.351), i.e., their abilities lie further below the average observed in the active workforce. As only the very worst workers are being replaced, the average difference between these mediocrities and their replacements is naturally larger (53% of a standard deviation).

The lower panel of Table 5 shows how these two counterfactual exercises would affect average ability in the entire sample of 22,697 worker-months for which we can estimate ability. In the counterfactual where we use the mediocrity threshold for rehires, the average ability of active workers would rise by 11% of the original standard deviation in estimated abilities. Using the mean novice as rehiring threshold also significantly increases the average ability of the active workforce (by 9.8% of a standard deviation), but this increase is significantly smaller than the one realized by moving to the mediocrity threshold.

< Insert Figure 4 around here >

In Figure 4, we plot the distribution of ability estimates for the actual and counterfactual workforce, where the level of observation is an active worker-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 workers of low to modest ability. The novice population, which replaces them, has a wider variance in ability, but on average outperforms the experienced workers they replace. In other words, firms face more uncertainty about the ability of these workers pre-hiring, but, as shown in the right tale of the distribution, the counterfactual novices possess upside potential, which is absent in the actual data. In panel B, we depict the same comparison for the full sample of observed worker-months in the data. These graphs illustrate that the counterfactual policy has a significant effect on the overall distribution of ability in the active labor force.

5.4. Which firms hire experienced mediocre workers instead of novices?

The theory we developed suggests that firms facing a difficult financial situation should be more inclined to hire mediocre workers instead of novices. Unfortunately, our data has no clear source of endogenous variation to test the causal effect of liquidity constraints. We can however assess whether mediocre hires are more prevalent at firms with less favorable finances. To this end we specify a linear probability model of the following form:

$$y_{mip} = \beta_x X_{ip} + \delta_{ip} + \eta_{mip} + \varepsilon_{mip}. \qquad (7)$$

In equation (7), y_{mip} is a variable which takes the value 1 if the hire of manager m by club i in period p was mediocre or substandard, and 0 if club i hired a novice instead. We do not consider hires of non-mediocre/substandard experienced workers, as the club should always prefer this option if it is feasible to attract a high ability tried-and-tested worker, regardless of its financial situation. Hence, the number of hires in the estimation sample will be smaller for the case of substandard hires, as less workers fall in this category. Our main explanatory variable of interest in X_{ip} is the club's wage-to-turnover ratio, which measures the club's total personnel cost divided by its revenues. This measure is a commonly used indicator of a club's financial position because wages are by far the most important cost to football clubs and often exceed the total annual revenues they earn (Peeters and Szymanski, 2014). We further control for the club's revenues and wage bill scaled by the industry average in the season the hire takes place. We finally add divisional and calendar month fixed effects (δ_{ip} and η_{mip}).

< Insert Table 6 around here >

The results of this analysis in Table 6 show that clubs with higher wage-to-turnover ratios tend to hire more mediocre and substandard workers instead of novices. Adding further controls to the regression does not substantially change this finding. As stated before, we cannot claim that these estimated effects have a causal interpretation. Interestingly, higher revenue clubs and clubs in higher tiers appear to be less inclined to hire novices. This further qualifies our finding above that clubs in lower divisions hire relatively more mediocre workers. Table 2 shows that lower division clubs are as likely to 'experiment' with novices as clubs in the top division, but they have less access to high-end experienced hires, which accounts for their propensity to hire mediocre and substandard experienced workers.

5.5. Are mediocre and substandard workers underperforming?

A potential concern with our analysis is that we may be classifying hires as mediocre or substandard who in retrospect turn out to be successful. This is particularly problematic if this misclassification is due to ability signals, which firms observe but we are unable to pick up. We examine this potential concern in Table 7. In the top panel we report the transition probabilities of rehired workers conditional on their rating as mediocre or non-mediocre at the start of their employment spell to being mediocre versus non-mediocre at the end of the spell. As in table 4, we report comparisons based on the score difference worker effects and the added win% measure. We find that the rating of a rehired worker is unlikely to change by the end of their employment spell, 82% of workers rated mediocre at the start of a spell were rated mediocre at the end, while 85% of those rated above the threshold were still rated above threshold by the end of their spell. A very similar pattern holds for the add win%-measure. This suggests that clubs have little reason for optimism that mediocre worker will get significantly better, and every reason to favor workers who are already rated above the mediocrity threshold. As shown in the bottom panel, we find that 96% of all non-substandard workers remain in this category by the end of their spell, whereas about 2 in 3 substandard experienced hires are still substandard by the end of their spell. This partly reflects that this threshold is less stringent, but it still clear that firms hiring non-substandard workers can safely expect to see sufficient worker performance.

< Insert Table 7 around here >

In appendix D, we follow Coupé et al (2006) in investigating whether past worker performance influences labor market outcomes as we would expect when employers act on this information. We find that workers with low estimated abilities are more likely to be fired than their more able counterparts. Upon losing their job, they are less likely to be rehired by another firm and less likely to start their next job in a more productive firm, where more money is at stake. These effects are stronger for low ability workers, who have more observable experience and whose ability can be inferred more precisely. We also examine the predictability of worker performance based on working experience. Figure D1 shows that ability estimates at the point of entry into the market have almost no predictive value, but within two months, past ability estimates explain upward of 80% of the variance in current ability estimates. For the remainder of the manager's career, the level of predictability is over 90%. In combination, these results confirm that employers need and use the information generated by past worker performance to formulate hiring policies.

6. Discussion and conclusion

In this paper, we have analyzed the role of liquidity constraints in the hiring of novice workers in a superstar labor market. We extend Terviö's (2009) model of hiring mediocrity to incorporate the effect of liquidity constraints. Under plausible assumptions about the firm's discount rate and the level of the liquidity constraints, firms retain experienced workers whose known ability is below the average ability of novice workers. Our model not only suggests a new reason for mediocre hiring, but the predicted inefficiency is also starker than in Terviö (2009). While rational from the firm's point-of-view, this "under-testing" of novice workers leads to a loss of social welfare in the sense that the average ability of the active labor force is below its long-run potential.

We empirically examine the hiring and re-hiring process of workers in a large dataset drawn from a high-stakes environment, English professional football, over a period of 38 years. The quality of the data allows us to introduce an explicit model of the timing at which worker ability is revealed. Since we know the exact date and result of each game played - the ultimate yardstick of worker performance - we can infer which information on worker performance "on-the-job" was available to firms at each point in time. Our estimation algorithm calculates updated worker effects à la Abowd et al (1999) after each calendar month in the dataset, taking newly revealed observations on worker performance into account. As such, we generate a career profile for each worker's estimated ability over time, rather than a single estimate based on their full career. We then investigate each instance where a firm hires an experienced worker, and compare the current ability estimate of this experienced worker to the average current ability estimates of recent labor market novices. In this way, we avoid hindsight in our analysis of firm hiring behavior.

Our analysis reveals that in more than forty percent of all instances where a firm hires an experienced worker, the experienced worker has an estimated ability below the “mediocrity” threshold, which would maximize the ability of the active labor force over time. The estimated ability for about one fifth of all rehired experienced workers is also below the mean ability of recently hired novices. A counterfactual exercise, which replaces these mediocre workers by hypothetical novice workers, shows that such a policy would significantly increase the ability of the active workforce over our sample period. As such, this paper provides the first, albeit indirect, empirical evidence of inefficient hiring in a real-life superstar labor market. We further found that lower division clubs with higher wage-to-turnover ratios are more inclined to choose a mediocre experienced worker over a novice. While we lack a source of exogenous variation to pin down causal effects in this analysis, this observation is in line with our theoretical model.

We have implicitly 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, e.g. allocating a fixed number of points among team in a sports league, 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. Moreover, leagues such as the English football leagues we analyze are engaged in competition with other leagues around the world to attract fans. In this sense, our findings suggest the existence of a market failure: a significant fraction of the rehires we observe are profit maximizing but reduce the average ability in the market below its full potential.

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; Palacios-Huerta, 2016), sports markets offer an outstanding laboratory for research on economic issues. In our dataset, the use of sports data enables us to observe all active workers in an entire industry, necessary for our analysis of turnover at the industry level. Moreover, we can credibly argue that testing novices within the industry is both costly and the only way to assess worker ability. As such, we exploit the richness of data available in sports to push beyond what is empirically feasible in other settings.

Acknowledgements

We thank the editor Raffaella Sadun and two anonymous referees for their insightful suggestions to improve this paper. We benefited from individual comments by Alex Bryson, Jan van Ours, Ignacio Palacios-Huerta, Michael Leeds and Arne Warnke. We are grateful to James Reade for sharing his Elo dataset for English football clubs. We thank seminar and conference participants at the University of Sankt Gallen, University of Michigan, Imperial College Business School, DIW Berlin, Erasmus School of Economics, Universitad Autonoma de Barcelona, Prague School of Economics, University of Reading, University of Zurich, 2016 Colloquium on Personnel Economics in Aachen, 2016 European Conference on Sports Economics in Groningen, 2017 European Association for Labour Economics meeting in Sankt Gallen, 2018 Southern Economic Association Meeting in Washington DC, 2018 Winter Meeting of Belgian Economists in Antwerp, and the International Football and Economics Conference in Baku.

Appendix A: Variable definitions

Name	Definition	Type
Table 1:		
• Games observed	Number of games worker is observed in the data, not taking into account potential missing variables.	Continuous
• Months observed	Distinct calendar months in which worker is observed in data, not taking into account potential missing variables.	Continuous
• Av. Age	Average age of worker in years over all his observations in data.	Continuous
• Av. Exp.	Average experience in games (both in England and abroad) over all his observations in data.	Continuous
• Av. Eng. Exp.	Average experience in games (confined to England) over all his observations in data.	Continuous
• Foreigner	Equals 1 if worker is not UK or Irish national, 0 otherwise.	Indicator
• Player-manager	Equals 1 if worker started career as player-manager, 0 otherwise.	Indicator
• Other man. exp.	Equals 1 if worker had other management function in football before 1st employment as manager, 0 otherwise.	Indicator
• Intern hire	Equals 1 if worker had other management function in same football club as 1st employment as manager, 0 otherwise.	Indicator
• Division	Division in which worker obtained first employment, ranked 4 (lowest) to 1 (highest).	Categorical
• Play prof.	Equals 1 if worker played as professional player, 0 otherwise.	Indicator
• Play big 4	Equals 1 if worker 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 worker has played for.	Continuous
• Ex-player club	Equals 1 if worker is employed by club that he also played for, 0 otherwise.	Indicator
• International	Equals 1 if worker 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
• Wage to revenue ratio	Total personnel costs incl. tax and social security payments divided by total turnover	Continuous
• Fixed assets	Total book value of fixed tangible assets of club over season, from financial accounts.	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
• # novice hires	Number of times club has hired novice worker over sample period.	Continuous
• Av. Tenure	Average tenure of managers at club 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
Table 6:		

• Wage to revenue ratio	Total personnel costs incl. tax and social security payments divided by total turnover	Continuous
• Relative revenues	Club revenues divided by the average club revenue in season	Continuous
• Relative wages	Club personnel cost divided by the average personnel cost in season.	Continuous
• Log hiring month number	Logarithm for a counter of the month of hiring within the data, August 1974=1, May 2011=380.	Continuous
Table D1:		
• Ability estimate	Ability estimate of worker 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. worker-firm pairing.	Continuous
• Log age	Logarithm of average age, expressed in days, in month of analysis.	Continuous
• Log experience	Logarithm of total experience, expressed in games, in month of analysis.	Continuous
Table D2:		
• Ability est. at spell end	Ability estimate of worker in month spell ends.	Continuous
• Average win%	Career average win percentage obtained by worker in month spell ends.	Continuous
• Average goals pro	Career average goals scored obtained by worker in month spell ends.	Continuous
• Log games	Logarithm of number of games worker has managed up until month spell ends.	Continuous
• Log age	Logarithm of age in days in month spell ends.	Continuous
• Foreigner	Equals 1 if worker 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

Appendix B: Results for alternative ability measures and comparison groups

In this appendix, we describe a set of alternative ways to gauge worker ability, which we implement in addition to the baseline measures reported in the main body of the paper. We consider two additional measures for on-field performance, which we use as alternative dependent variables to estimate the worker effects in equation (5). First, we replace the score difference with the win% per game, recording a loss as 0, a draw as 0.5 and a win as 1. Second, we consider the update in the club's Elo points after the game. The Elo points reflect the quality of a club based on past sporting results. After a win (loss) the Elo value is updated upward (downward) by an amount, which depends on the pre-game difference in Elo scores between the competing clubs. As shown by Hvattum and Arntzen (2010), the pre-game Elo points are an excellent predictor of game outcome.

We experiment with several alternative estimation methods for the baseline model in equation (5). We first address potential feedback effects from game results to input values, 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 '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 worker effects using additional controls for worker experience, i.e. the level, square and cube of log career games managed ('exp.' below). Third, we allow for the possibility that observations that are more recent provide more information on current worker 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 'weight' 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, Lazear et al., 2015 and Peeters et al., 2020 for more). In this approach, referred to as 'Spell' below, we estimate firm-worker spell effects in the first stage, and then split these into worker 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 worker has over his career in the data but allows for differing match qualities among spells.

Finally, we conduct two robustness checks regarding to the sample we use for our analysis. In the "+20 obs." subsample we focus on worker hires for whom the worker effect is estimated on at least 20 game observations. This should reduce the noise in the analysis, which may arise from the rehiring of managers with very short initial spells. We also vary the group of novices to which we compare the experienced hires. In our baseline we look at all first-time workers from the past ten years. Here we reduce that comparison period to five years.

< Insert Table B1 and Table B2 around here >

Table B1 shows the summary statistics of the alternative ability measures we estimate, both in terms of game results measures and estimation methods. In Table B2 we report the correlation between the baseline measures from the main body of the paper 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. The two alternative game result measures (win% and Elo) also correlate highly with both worker performance measures reported in the main body of the paper.

< Insert Table B3 around here >

As in Table 4 in the main body of the paper, the top panel of Table B3 shows the analysis for the mediocre threshold for the alternative ability measures and comparison groups, whereas the bottom panel uses the substandard threshold. The robustness checks with Elo and win% largely confirm our main results on hiring. If anything, using Elo points leads to a higher estimate of the share of mediocre hiring. This is because using the Elo points favors recently entered managers vis-à-vis more experienced managers. As van Ours and van Tuijl (2016) show, soccer managers generally get replaced following a string of bad results. A bad streak of results lowers the club's Elo rating and makes it easier to improve over the next set of games, a phenomenon van Ours and van Tuijl ascribe to "reversion to the mean". In terms of our estimates of worker effects, recent entrants are proportionally more impacted by this early reversion to the mean than the estimates of their more experienced counterparts. This effect is absent when using either the goal difference or game outcome as the dependent variable for the worker ability regressions. As such, it makes sense that we find more 'mediocre' hiring using the Elo model than other output measures. Using the alternative estimation methods changes very little in the analysis of mediocre hiring. This is not very surprising given the very high correlation between these measures and our baseline measure. All results point in the same direction, i.e. between around 25% and 45% of all hiring events involve a mediocre experienced worker. For the average novice threshold lies between around 21% and 29%.

Appendix C: Connectedness test for AKM regression model

In this appendix, we look into the connectedness of the underlying worker-firm network in our dataset. As shown by Jochmans and Weidner (2019), the estimation error in the worker fixed effects critically depends on the density of the connections in this network. When the connected network is too sparse, the implied estimation error of the fixed effects explodes, which may bias any higher-order moments calculated from these estimates. This problem is in effect a generalization of the "limited mobility bias" discussed by Andrews et al. (2008) in relation to assortative matching in the labor market. In our case, this problem could for example affect the estimated variance of the novice distribution, which we base on the estimated novice fixed effects.

We test for the severity of this issue through the approach of Jochmans and Weidner (2019) to calculate the implied bias in the variance of the worker fixed effects. Adopting the notation in Jochmans and Weidner (2019), we first construct the adjacency matrix A_p for each of the 280 worker-firm networks on which the estimation for each month p is based. We weigh the importance of each connection (i.e. each moving worker) by the number of observations we have for the worker at each firm in the worker's career. We then calculate the Laplacian $\hat{L}_p^*$ and normalized Laplacian S_p of A_p. These matrices characterize the connectedness of the worker-firm network for each sample in the moving estimation algorithm. Therefore the first non-zero Eigenvalue of S_p, λ_{2p} should be "bounded away" from 0 for the network to be sufficiently connected. In Figure C1, we plot these values for each monthly sample in our rolling algorithm. The eigenvalues increase as our sample progresses, which is intuitive because the number of club effects grows at a slower pace than the connections between the clubs. Apart from the very start of the rolling estimation, the eigenvalues are (much) larger than 0 and clearly above the 0.0039 reported by Jochmans and Weidner (2019) as an example of a weakly connected network. Hence, the worker-firm networks in our analysis are sufficiently connected.

< Insert Figure C1 around here >

Using the matrices defined above, we can calculate the implied bias of the plug-in estimator of the variance of the worker fixed effects measured as a percentage of the error variance. We show the results on the right-hand axis of Figure C1. In line with our finding for the connectedness, we find that the implied bias declines sharply as we move away from the very earliest subsamples. By the 20th monthly subsample of our rolling estimation (month 120 of the dataset), the potential bias stays below 1% for the remainder of the analysis. By way of comparison, the example of a well identified network given in Jochmans and Weidner (2019) yields an implied bias of 5.8%, much higher than the levels we uncover here.

Appendix D: Supporting Empirical Results

1. Worker ability and attrition

If firms are able to observe the ability of their own worker with increasing accuracy over time, we should expect them to terminate their 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

$$y_{mip} = \beta_q \hat{q}_{mp} + \beta_{dq} \Delta\hat{q}_{mp-1} + \beta_x X_{mp} + \varepsilon_{mp}. \qquad (D1)$$

The dependent variable, y_{mip}, is an indicator, which equals 1 if the worker's tenure at team i ended in month p.[16] We explain spell termination by two main variables, (a) the current ability estimate of the worker, $\hat{q}_{mp}$, 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 worker characteristics containing log of age, tenure and experience and add FEs for the year, calendar month and division.

< Insert Table D1 around here >

Table D1 confirms that workers 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 of around 2%, while the baseline termination 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 workers, and even more likely to retain them when they are on an upward trajectory.

2. Worker ability and career dynamics

Finally, we assess the observability of worker ability to rival firms. First, we examine whether higher estimated ability increases the probability of survival in the labor market. 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 estimate of worker ability in month p, and a vector of personal characteristics, X_{mp}, as explanatory variables, i.e.

$$y_{mp} = \beta_q \hat{q}_{mp} + \beta_x X_{mp} + \varepsilon_{mp}. \qquad (D2)$$

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.

< Insert Table D2 around here >

Table D2 reports the model estimates for our sample. The model variants display a strikingly consistent pattern: workers 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 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, younger and foreign workers. 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).

3. Predicting worker ability at entry and speed of convergence

As can be seen in Table 1, a good deal is known about most workers at the beginning of their career, even when 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 worker ability. Moreover, for our model to make sense, future ability estimates have to be predictable after a worker has entered the labor market. We examine these assumptions by regressing the current worker ability estimate on the ability estimate (a) in the previous month and (b) five months earlier. Thus our forecasting model for novice ability is,

$$\hat{q}_{mp+k} = \beta_x X_{mp} + \beta_q \hat{q}_{mp} + \varepsilon_{mp}. \qquad (D3)$$

Here, $\hat{q}$ is the ability estimate and k is the number of time periods we predict ahead (one or five). The vector X_{mp} includes worker characteristics in month p, in particular, log 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 entry, we also run this regression excluding information we observe on the job, i.e.

$$\hat{q}_{m,p+k} = \beta_x X_{m,p} + \varepsilon_{mp}. \qquad (D4)$$

We focus on the adjusted R^2 of these regressions to identify the speed of convergence. The graphs of the results for 1 and 5 months ahead are shown in Figure D1. 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 months ahead model the adjusted 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. These results suggest that ability is revealed relatively fast in this labor market. By comparison, Lange (2007) reports that the initial expectation errors of employers fall by 50% within 3 years. The adjusted R^2 of the model based on pre-entry characteristics never moves away from zero, which implies it is hard to predict the ability of novices from observable characteristics prior to their entry in the labor market. Note that the entry mode into the labor market also has very limited predictive power for worker ability. In other words, firms are not very successful in selecting managers based on the private information they may have obtained during the manager's previous employment in other roles at the club.

< Insert Figure D1 around here >

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

Endnotes

[1] Corresponding author, e-mail: peeters@ese.eur.nl, postal: Erasmus School of Economics, Postbus 1738, 3000 DR Rotterdam, Netherlands, phone: +32 494 12 49 36

[2] E-mail: stefansz@umich.edu, postal: 1402 Washington Hts., Ann Arbor, MI 48109-2013, United States.

[3] E-mail: marko.tervio@aalto.fi, postal: PO BOX 21210, FI-00076 Aalto, Finland

The authors agree to make all relevant datasets and estimation codes publicly available in a permanent online repository.

[4] See e.g. Cohen et al. (2018) and Peeters et al (2020) for analyses of managerial discretion and performance using data from US basketball and baseball.

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

[6] In recent years some high-profile managers have obtained multiple (e.g. four) year deals. This is a very recent evolution and it has never been and still is not a standard practice for novice managers.

[7] Apart from a brief expansion in the early 1990s, this number has been constant. Due to promotion and relegation to and from the semi-professional (regional) level, there are more than 92 unique clubs in the dataset.

[8] The exchange rate between the US dollar and the UK pound has fluctuated between $1.08 and $2.12 to the pound over the period considered in this paper, averaging $1.63.

[9] 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 similar to the analysis of Bertrand and Schoar (2003). See Muehlheusser et al. (2018) for an application of this method to football managers in the German Bundesliga.

[10] Within season manager turnover is high, such that choosing a longer updating period, e.g. a season, could result in large time lags between a hiring/firing and the last update to the manager ability estimates.

[11] We use Graham Taylor as reference manager, as he has a long career and is present from the start of the sample.

[12] We vary this timeframe to assess the robustness of our findings, see appendix B.

[13] Our sample of novices is by default restricted to workers who were hired at some point. Hence, we implicitly assume that the available novice workers, who were not hired, are equivalent to novice workers who were selected. If firms would have effective screening mechanisms for potential novices, our approach may overestimate the ability of unselected novices. However, in appendix D we show that worker characteristics, which are observable before labor market entry, do not predict the ability of novice workers very well. If successful screening exists, it should largely be based on unobservable worker treats, something which we unfortunately cannot verify. While the pool of potential novices is in principle global, we only consider workers who start their careers in England.

Foreign managers who come to England are relatively scarce (Table 1) and tend to be at the upper end of the ability distribution. Since we do not have data on the performance of novices abroad, we cannot say how representative these migrant managers are of all potential foreign novices. Excluding these foreign novices means that we may understate the expected ability of the true (global) novice pool. Note that a worker coming in after a career abroad is neither 'rehired', nor part of the novice distribution. A manager coming to England with solely foreign experience enters our analysis only if he obtains a second employment spell in England.

[14] For example, in Figure 1, McMahon is estimated to be mediocre at the start of his second employment spell, therefore only his first spell is included in the counterfactual data of manager-months.

[15] 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 AKM estimation is not fully taken into account. However, bootstrapping the entire rolling estimation procedure would take up a prohibitive amount of computer time.

[16] 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.

References

Abowd, J. M., Kramarz, F., & Margolis, D. M. (1999). High Wage Workers and High Wage Firms. *Econometrica*, *67*(2), 251-333. doi:10.1111/1468-0262.00020

Andrews, M. L., Gill, L., Schank, T., & Upward, R. (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society A*, *171*(3), 673-697. doi:10.1111/j.1467-985X.2007.00533.x

Bertrand, M., & Schoar, A. (2003). Managing with Style: the Effect of Managers on Firm Policies. *Quarterly Journal of Economics, CXVIII*(4), 1169-1208. doi:10.1162/003355303322552775

Bridgewater, S., Kahn, L. M., & Goodall, A. H. (2011). Substitution and complementarity between managers and subordinates: Evidence from British football. *Labour Economics*, *18*, 275-286. doi:10.1016/j.labeco.2010.10.001

Caggese, A., Cunat, V., & Metzger, D. (2019). Firing the wrong workers: Financing constraints and labor misallocation. *Journal of Financial Economics, 133*(3), 589-607. doi:10.1016/j.jfineco.2017.10.008

Cohen, A., Levy, N., & Sasson, R. (2018). Termination Risk and Agency Problems: Evidence from the NBA. *Journal of Law, Economics and Organization*, *34*(4), 579-618. doi:10.1093/jleo/ewy018

Cornelissen, T. (2008). The Stata command felsdvreg to fit a linear model with two high–dimensional fixed effects. *Stata Journal, 8*(2), 170-189. doi:10.1177/1536867X0800800202

Coupé, T., Smeets, V., & Warzynski, F. (2006). Incentives, Sorting and Productivity along the Career: Evidence from a Sample of Top Economists. *Journal of Law, Economics and Organization, 22*(1), 137-167. doi:10.1093/jleo/ewj010

Duggan, M., & Levitt, S. D. (2002). Winning Isn't Everything: Corruption in Sumo Wrestling. *American Economic Review*, *92*(5), 1594-1605. doi:10.1257/000282802762024665

Farber, H. S., & Gibbons, R. (1996). Learning and Wage Dynamics. *Quarterly Journal of Economics, 111*(4), 1007–1047. doi:10.2307/2946706

Garicano, L., Palacios-Huerta, I., & Prendergast, C. (2005). Favoritism Under Social Pressure. *Review of Economics and Statistics, 87*(2), 208–216. doi:10.1162/0034653053970267

Hope, C. (2003). When should you sack a football manager? Results from a simple model applied to the English Premiership. *Journal of the Operational Research Society, 54*(11), 1167-1176. doi:10.1057/palgrave.jors.2601621

Hvattum, L. M., & Arntzen, H. (2010). Using ELO ratings for match result prediction in association football. *International Journal of Forecasting, 26*, 460-470. doi:10.1016/j.ijforecast.2009.10.002

Jackson, K. C. (2013). Match Quality, Worker Productivity, and Worker Mobility: Direct Evidence from Teachers. *The Review of Economics and Statistics, 95*(4), 1096–1116. doi:10.1162/REST_a_00339

Jochmans, K., & Weidner, M. (2019). Fixed Effects Regressions on Network Data. *Econometrica*, *87*, 1543-1560. doi:10.3982/ECTA14605

Jovanovic, B. (1979). Job matching and the theory of turnover. *Journal of Political Economy, 87*(5), 972-990. doi:10.1086/260808

Kahn, L. M. (2000). The Sports Business as a Labor Market Laboratory. *Journal of Economic Perspectives, 14*(3), 75-94. doi:10.1257/jep.14.3.75

Lange, F. (2007). The Speed of Employer Learning. *Journal of Labor Economics, 25*(1), 1-35. doi:10.1086/508730

Lazear, E. P., Shaw, K. L., & Stanton, C. (2015). The Value of Bosses. *Journal of Labor Economics, 33*(4), 823-861. doi:10.1086/681097

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

Muehlheusser, G., Schneemann, S., Sliwka, D., & Wallmeier, N. (2018). The Contribution of Managers to Organizational Success: Evidence from German Soccer. *Journal of Sports Economics*, *19*(6), 786–819. doi:10.1177/1527002516674760

Palacios-Huerta, I. (2016). *Beautiful game theory: How soccer can help economics*. Princeton: Princeton University Press.

Pallais, A. (2014). Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review, 104*(11), 3565–3599. doi:10.1257/aer.104.11.3565

Peeters, T., & Szymanski, S. (2014). Financial Fair Play in European football. *Economic Policy, 29*(78), 343-360. doi:10.1111/1468-0327.12031

Peeters, T., Salaga, S., & Juravich, M. (2020). Matching and Winning? The Impact of Upper- and Middle Managers on Team Performance in Major League Baseball. *Management Science, 66*(6), 2735-3751. doi:10.1287/mnsc.2019.3323

Stanton, C., & Thomas, C. (2016). Landing the First Job: The Value of Intermediaries in Online Hiring. *Review of Economic Studies, 83*(2), 810-854. doi:10.1093/restud/rdv042

Sterling, A. D. (2014). Friendships and Search Behavior in Labor Markets. *Management Science, 60*(9), 2341–2354. doi:10.1287/mnsc.2013.1857

Strobl, G., & Van Wesep, E. D. (2013). Publicizing Performance. *Management Science, 59*(4), 918-932. doi:10.1287/mnsc.1120.1620

Szymanski, S. (2017). Entry into exit: insolvency in English professional soccer. *Scottish Journal of Political Economy, 64*(4), 419-444. doi:10.1111/sjpe.12134

Szymanski, S., & Smith, R. (1997). The English football industry: profit, performance and industrial structure. *International Review of Applied Economics, 11*(1), 135-153. doi:10.1080/02692179700000008

Terviö, M. (2009). Superstars and Mediocrities: Market Failure in the Discovery of Talent. *Review of Economic Studies, 76*, 829–850. doi:10.1111/j.1467-937X.2008.00522.x

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), 591-604. doi:10.1111/ecin.12280

Figure captions

Figure 1: Optimal retention threshold ψ as function of liquidity constraint c discount factor δ

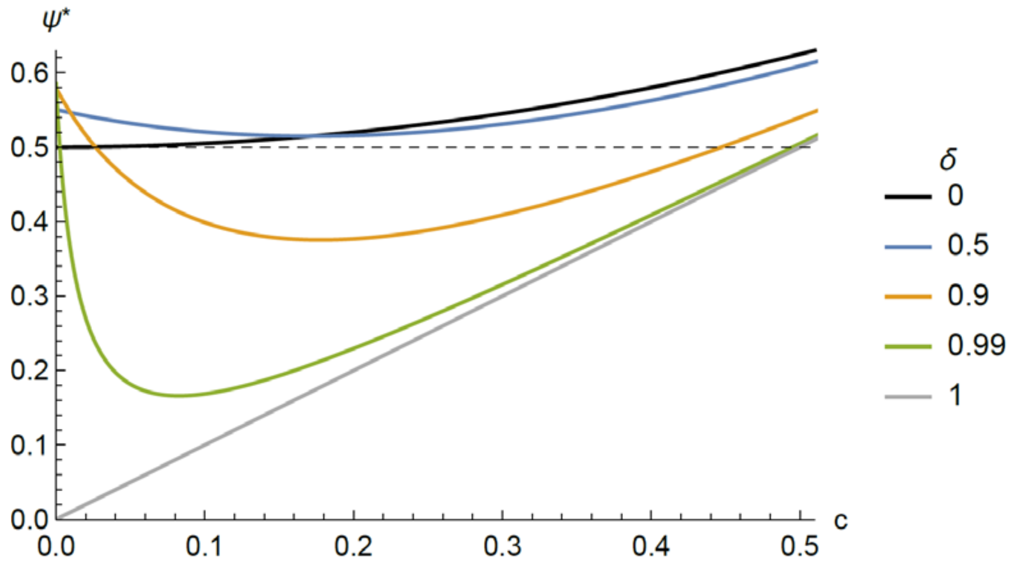

Notes: The lines show the firm's optimal retention thresholds as a function of the level of the liquidity constraint (c) on the x-axis and the discount factor (δ). We derive these from the theoretical model we sketch in section 2 with uniformly distributed worker ability. Cases below the dashed line (population mean) indicate "substandard" hiring. The limiting case of infinite patience (delta approaches 0) results in the maximization of average steady state ability.

Figure 2: Career evolution of ability estimates for selected workers

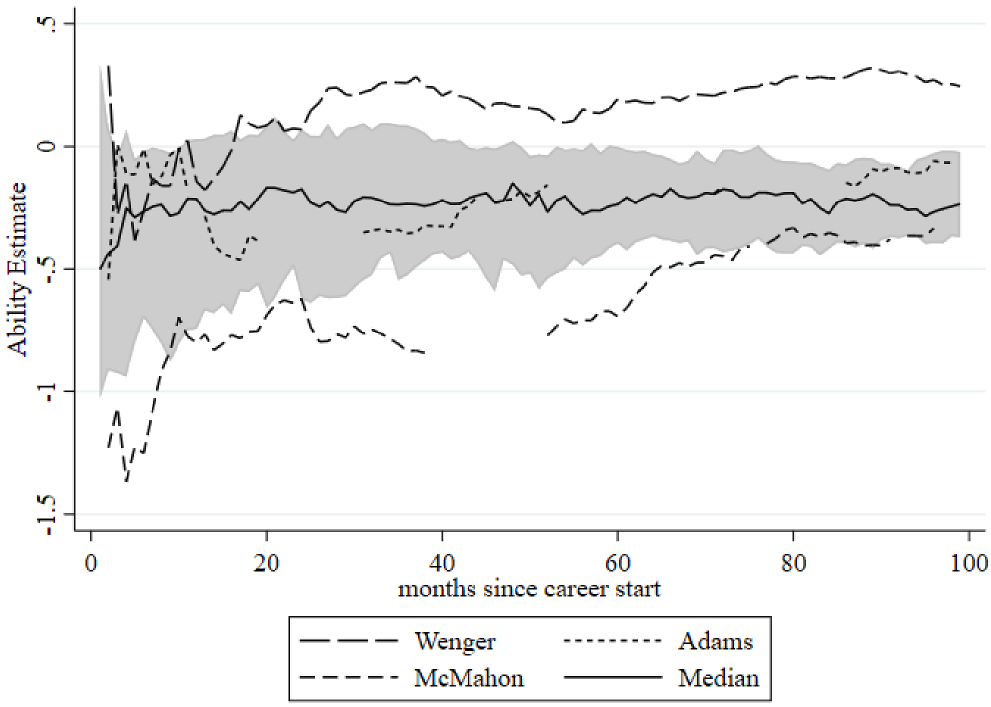

Notes: The plots depict the evolution of the ability estimates of Arsène Wenger, Micky Adams and Steve McMahon over the first ten years of their careers. The shaded area is the interquartile range for the ability estimates of managers with the same experience in number of months since labor market entry.

Figure 3: Ability estimates of rehired experienced workers versus mean of recent novice workers

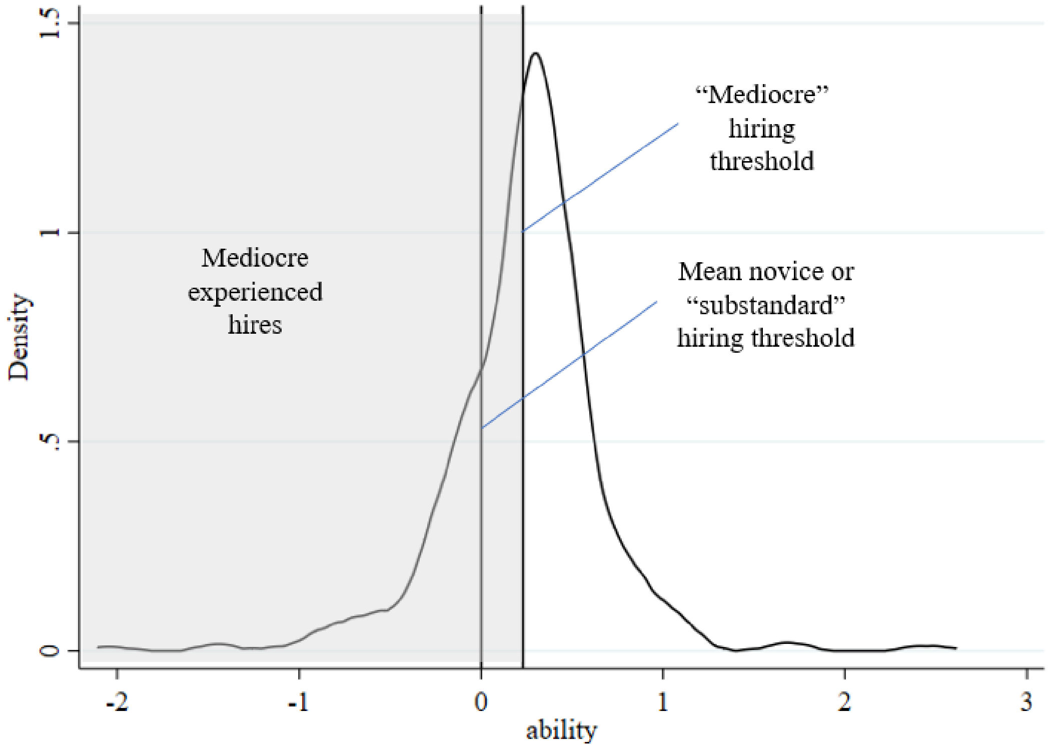

<u>Notes</u>: The figure depicts the distribution of the ability estimates of rehired workers vis-à-vis the average novice worker. We graph the Kernel density for the estimated worker ability minus the mean novice ability at the time of hiring. We show results using a 10-year time window of past novices.

Figure 4: Distribution of worker ability estimates in actual and counterfactual workforce

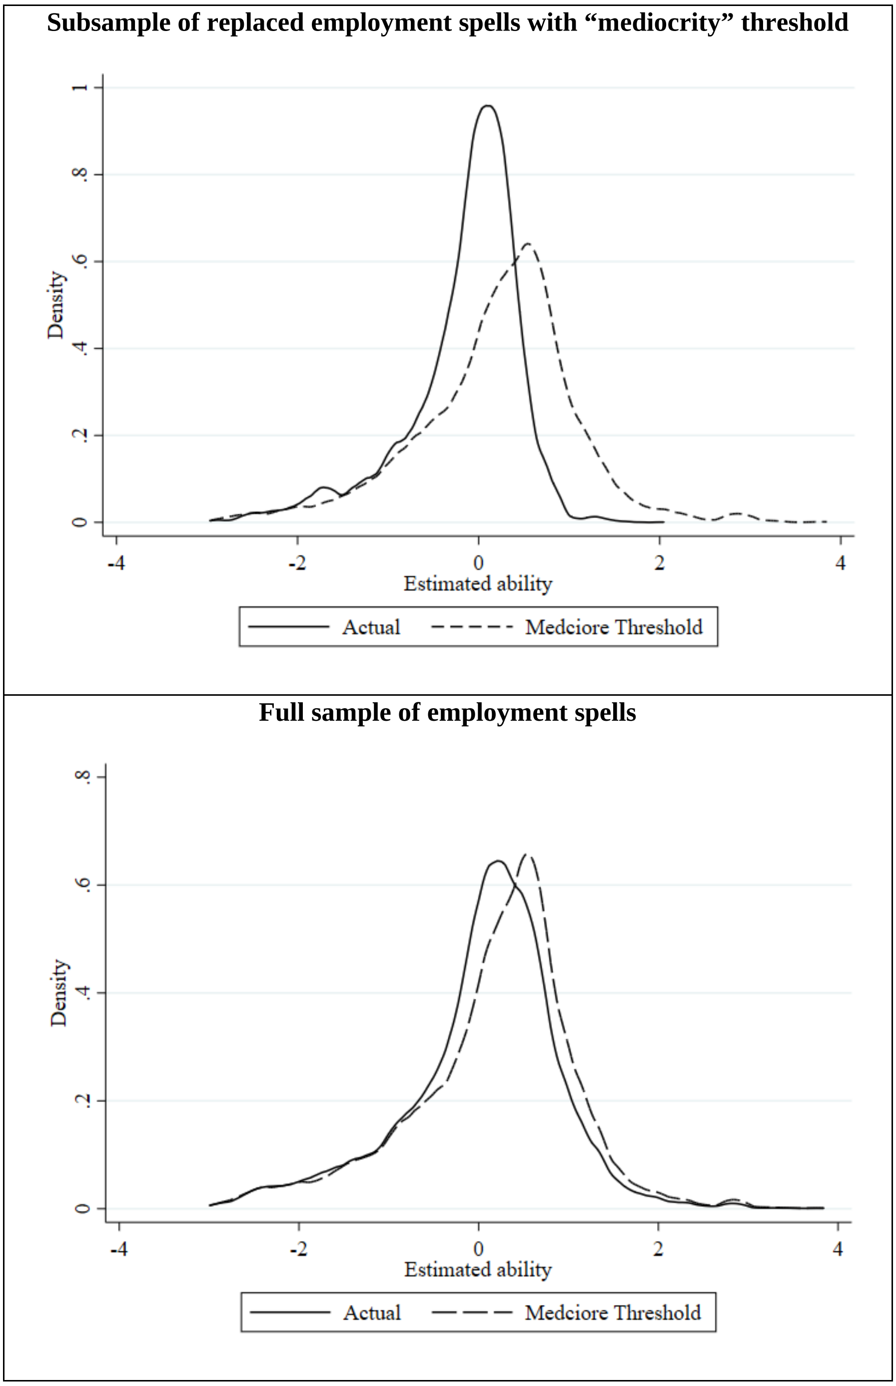

Notes: The graphs depict the distribution of contemporary ability estimates in the actual and counterfactual worker populations. The top panel focuses on the replaced subsample, i.e. end career spells of mediocre workers. The bottom panel depicts the full sample results. The level of observation is a worker-month, so each worker 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 C1: Global connectedness and implied bias in the variance of worker FEs for each monthly estimation sample

2nd Eigenvalue of normalized Laplacian for adjacency matrix over sample periods

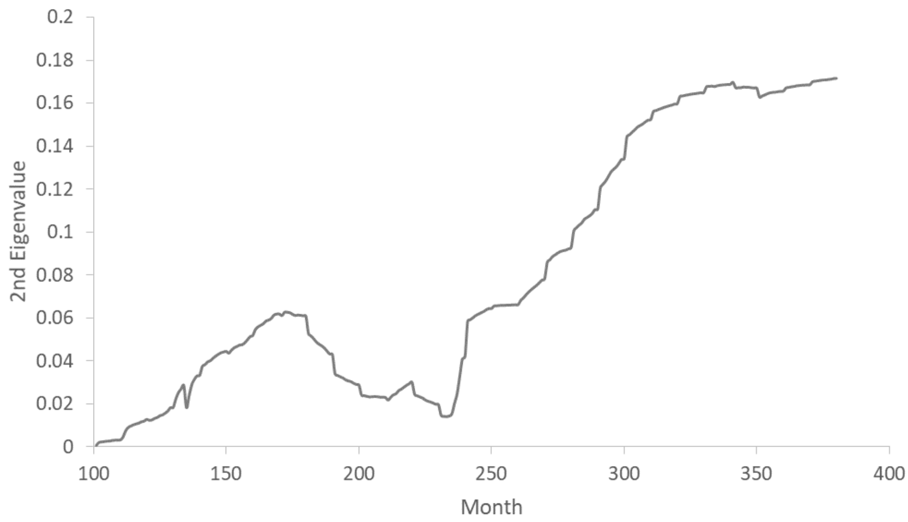

Implied bias in variance of worker effects as % of error variance

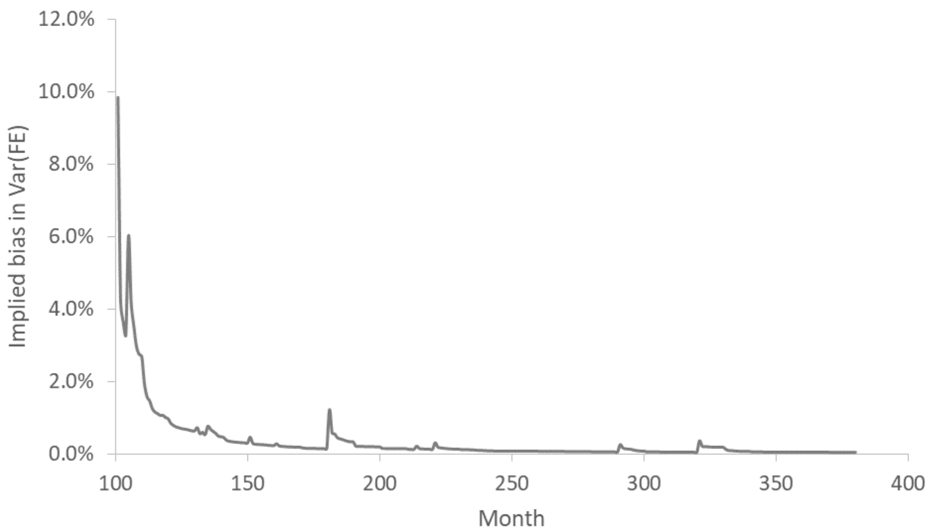

Notes: The top panel of this figure shows the 2nd eigenvalue of the normalized Laplacian of the adjacency matrix for each subsample in the rolling AKM estimation algorithm. The bottom panel depicts the implied bias in the variance of the worker fixed effects as a % of the error variance. The ‘month’ values on the x-axis refer to the last month included in the subsample.

Figure D1: Adjusted R-squared for regression of worker ability estimates 1 and 5 months ahead on entry characteristics and current ability estimates, by number of months observed

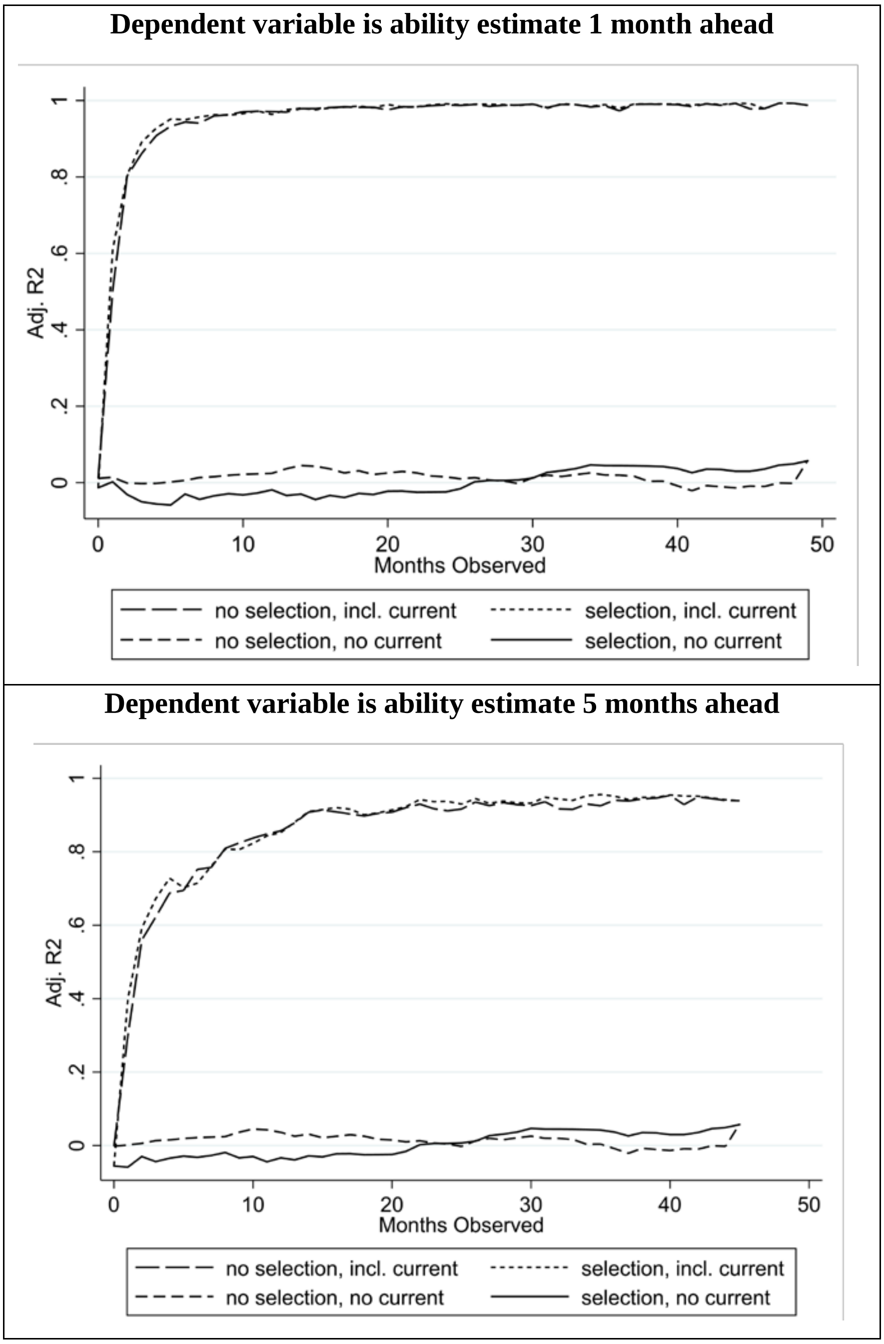

<u>Notes</u>: To form these graphs we draw subsamples of the dataset by the number of months a worker has been observed in the data. In each subsample we regress worker ability estimates 1 month ahead (top panel) and 5 months ahead (lower panel) on (a) a set of worker characteristics at labor market 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 workers ('no selection') and for the subset of workers, that attains at least 50 monthly observations ('selection').

Table 1: Summary statistics at individual worker 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	201.3	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
Worker Characteristics											
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	153.0	114.9	113.5	272.1	152.0
Av. Eng. Exp. (game)	830	115.6	141.2	27.2	98.8	49.1	101.3	88.3	76.8	257.5	139.8
Foreigner	830	0.060	0.238	0.098	0.298	0.078	0.270	0.060	0.239	0.021	0.145
Market Entry											
Player-manager	827	0.295	0.456	0.244	0.431	0.240	0.428	0.297	0.458	0.374	0.485
Other man. exp.	827	0.518	0.500	0.580	0.495	0.555	0.498	0.500	0.501	0.464	0.500
Intern hire	827	0.335	0.472	0.458	0.500	0.362	0.482	0.322	0.467	0.251	0.435
Division	942	2.785	0.982	2.807	1.050	2.978	0.999	2.815	0.973	2.543	0.849
Playing history											
Play prof.	830	0.954	0.209	0.910	0.288	0.952	0.214	0.961	0.194	0.975	0.158
Play big 4	830	0.665	0.472	0.602	0.491	0.683	0.466	0.625	0.485	0.723	0.448
Num. Eng. Team	830	3.463	2.291	3.290	2.631	3.617	2.353	3.362	2.298	3.509	2.003
Ex-player club	830	0.435	0.448	0.530	0.499	0.462	0.489	0.464	0.459	0.326	0.330
International	830	0.375	0.484	0.308	0.464	0.339	0.474	0.431	0.496	0.391	0.489

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

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.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Financial data (2010 k£)											
Years observed	98	28.9	8.85	20.9	9.54	29.5	9.00	31.2	6.11	33.8	4.18
Wages	98	4763	6999	510	460	1252	715	4119	2761	12974	9400
Revenue	98	7339	12352	707	712	1787	937	5424	3847	21099	18003
Wage to turnover ratio	98	70.1%	11.7%	73.7%	8.0%	72.4%	9.4%	73.4%	13.1%	61.1%	11.2%
Fixed assets	98	10764	23603	662	643	1471	972	5948	4812	34383	37838
Sports results											
Goals pro	98	57.8	3.60	58.0	3.48	59.2	2.60	57.1	3.04	57.0	4.70
Goals against	98	57.5	5.73	62.8	3.56	59.0	1.95	57.2	2.90	51.4	6.26
Goal Difference	98	0.28	7.19	-4.86	5.33	0.21	3.11	-0.09	3.24	5.64	10.26
Points	98	57.3	4.70	54.9	4.78	58.9	3.97	57.3	3.57	57.9	5.50
Win %	98	50.2%	3.8%	47.8%	2.6%	50.1%	1.9%	49.9%	2.0%	53.0%	5.6%
Managerial hiring											
Number hires	98	17.1	6.89	13.0	6.34	18.2	6.97	19.0	5.75	17.9	7.00
# novice hires	98	7.93	3.92	8.00	4.02	9.60	4.01	8.46	3.86	5.68	2.79
Av. tenure	98	182	95.9	177	90.6	163	72.7	165	45.8	221	140
Av. hires/year obs.	98	0.60	0.18	0.64	0.17	0.64	0.19	0.61	0.15	0.53	0.20

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 worker level by quartile in estimated ability

		Full sample		Ability quart 1		Ability quart 2		Ability quart 3		Ability quart 4	
	Obs.	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Score difference worker effects											
Mean value	626	-0.302	1.282	-1.947	1.157	-0.463	0.237	0.160	0.170	1.051	0.605
Number monthly estimates	626	36.3	43.3	16.5	20.3	34.4	39.4	55.5	51.9	38.7	45.7
Total observations in data	626	141.4	175.5	89.2	114.5	139.2	169.2	199.9	210.2	137.2	177.8
Win% vs. prior club average											
Mean value	752	-0.476	1.603	-2.514	1.297	-0.672	0.228	0.033	0.210	1.245	1.036
Number monthly estimates	752	33.8	42.7	9.660	12.6	31.6	31.0	51.7	48.9	42.3	53.2
Total observations in data	752	158.3	187.3	44.7	55.8	147.5	134.3	251.0	216.4	189.9	224.1

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

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

Ability threshold	Ability Measure	Division	# novices compared	# Spells	# Spells < threshold	% Spells < threshold	# Rehired workers	# workers < threshold	% workers < threshold
Mediocre	Score dif. effect	All	201	755	326	43.2%	322	192	59.6%
		1	201	142	47	33.1%			
		2	201	235	90	38.3%			
		3	200	181	84	46.4%			
		4	201	197	105	53.3%			
	Added win%	All	255	819	333	40.7%	355	203	57.2%
		1	250	150	53	35.3%			
		2	250	250	91	36.4%			
		3	251	202	90	44.6%			
		4	250	217	99	45.6%			
Substandard	Score dif. effect	All	201	755	171	22.6%	322	123	38.2%
		1	201	142	22	15.5%			
		2	201	235	40	17.0%			
		3	200	181	50	27.6%			
		4	201	197	59	29.9%			
	Added win%	All	255	819	129	15.8%	355	100	28.2%
		1	250	150	20	13.3%			
		2	250	250	30	12.0%			
		3	251	202	36	17.8%			
		4	250	217	43	19.8%			

Notes: We depict results for the score difference worker effects and the added win% compared to previous managers. Columns 5-8 show the absolute and relative number of employment spells where an experienced worker was mediocre at the time of hiring. Columns 9-11 show the absolute and relative number of individual workers with multiple employment spells who started at least one employment spell as a mediocre hire. In the top panel we show results for the mediocrity threshold, the bottom panel uses the mean novice as hiring threshold. Alternative ability measures and assumptions on the novice comparison group lead to comparable results. See appendix A for more detail.

Table 5: Summary ability and career length for counterfactual workforce

	# Worker-months	Mean contemporary ability estimate	Bootstrap std. error
Replaced spells mediocrity threshold			
Actual ability	7,096	-0.1752	0.0083
Counterfactual ability	7,096	0.1791	0.0274
Average ability difference	7,096	0.3543	0.0296
Replaced spells substandard threshold			
Actual ability	4,227	-0.3509	0.0112
Counterfactual ability	4,227	0.1766	0.0362
Average ability difference	4,227	0.5274	0.0387
Full sample ability difference			
Actual vs. mediocrity threshold	22,697	0.1108	0.0094
Actual vs. substandard threshold	22,697	0.0982	0.0074
Substandard vs. mediocrity threshold	22,697	0.0126	0.0047

Notes: This table shows the results of a counterfactual, where we drop all employment spells in worker careers, which occur after a worker has been hired as a mediocre/substandard worker. We replace these spells by draws from the observed novice ability distribution, again disregarding spells following a mediocre/substandard hire. We show the average difference in estimated ability for the subsample we replace (top two panels) and for the full sample (bottom panel). We normalize the ability estimates by subtracting the overall average and dividing by the standard deviation such that these numbers measure standard deviations in ability. The level of observation is the individual worker-month, with the number of observations indicated in the first column. We perform 100 bootstrap replications of this counterfactual to assess the significance of the difference between the actual and counterfactual samples.

Table 6: LPM results for probability that firm hires novice worker instead of 'mediocre' or 'substandard' experienced worker

Firm hires:	Mediocre worker (1) vs. Novice (0)			Substandard worker (1) vs. Novice (0)		
Wage to revenue ratio	0.153**	0.240***	0.174**	0.157**	0.238**	0.169*
	(0.070)	(0.083)	(0.083)	(0.079)	(0.096)	(0.097)
Relative revenues		0.119**	0.086*		0.126*	0.071
		(0.053)	(0.052)		(0.069)	(0.069)
Relative wages		-0.040	-0.011		-0.033	0.004
		(0.056)	(0.055)		(0.068)	(0.068)
Log hiring month number			0.117***			0.081**
			(0.034)			(0.038)
Division 1		Baseline			Baseline	
Division 2	-0.149***	-0.037	-0.035	-0.150***	-0.020	-0.037
	(0.048)	(0.064)	(0.063)	(0.055)	(0.072)	(0.071)
Division 3	-0.205***	-0.066	-0.066	-0.210***	-0.046	-0.068
	(0.048)	(0.074)	(0.073)	(0.054)	(0.082)	(0.081)
Division 4	-0.242***	-0.095	-0.080	-0.214***	-0.041	-0.050
	(0.049)	(0.076)	(0.076)	(0.055)	(0.085)	(0.084)
Constant	-0.149***	-0.037	-0.035	-0.150***	-0.020	-0.037
	(0.048)	(0.064)	(0.063)	(0.055)	(0.072)	(0.071)
Calendar month FE	No	No	Yes	No	No	Yes
Observations	990	990	990	773	773	773
R-squared	0.028	0.038	0.084	0.025	0.036	0.087

Notes: Table reports regression results for a linear probability model where the dependent is an indicator equaling 1 if a rehire is mediocre/substandard, and 0 if hire is a novice. We do not report point estimates for the hiring month FE to aid readability. Standard errors are given in parentheses, *** denotes significance at 1% level, ** at 5% level, * at 10% level.

Table 7: Transition probabilities from 'mediocre' to 'non-mediocre' during employment spell

Mediocrity threshold		End of Spell: non-mediocre	End of Spell: mediocre	Observations	Ability Estimator
Start of Spell	non-mediocre	85.1%	14.9%	429	Score dif. FE
	mediocre	17.5%	82.5%	326	
	non-mediocre	86.8%	13.2%	486	Add win%
	mediocre	18.9%	81.1%	333	
Substandard threshold		**End of Spell: non-substandard**	**End of Spell: substandard**	**Observations**	**Ability Estimator**
Start of Spell	non-substandard	95.7%	3.3%	584	Score dif. FE
	substandard	33.9%	66.1%	171	
	non-substandard	96.3%	3.7%	690	Add win%
	substandard	38.0%	62%	129	

Notes: Table displays the transition probabilities for a worker to move from mediocre (substandard) to non-mediocre (non-substandard) over the course of an employment spell. At both the start and end of the employment spell, we compare the ability estimate of the experienced worker to the contemporaneous novice distribution. Hence, both movements in the pool of available entrants and updates to the worker ability estimates may cause shifts in the categorization.

Table B1: Summary statistics alternative ability estimation methods

	Full sample			Ability quart 1		Ability quart 2		Ability quart 3		Ability quart 4	
Estimation method	**Obs.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Alternative outcome measures											
Win% worker effect	626	-0.275	1.298	-1.969	1.002	-0.471	0.233	0.204	0.175	1.145	0.698
Elo worker effect	626	-0.078	1.493	-1.814	1.304	-0.291	0.179	0.266	0.158	1.535	1.173
Score difference worker effects											
IV for wage bill	626	-0.310	1.269	-1.924	1.081	-0.535	0.234	0.155	0.176	1.071	0.638
Experience polynomial	626	-0.370	1.294	-2.012	1.168	-0.539	0.205	0.065	0.166	1.015	0.656
1st stage spell effect	626	-0.274	1.206	-1.832	0.996	-0.469	0.234	0.188	0.155	1.024	0.618
Weighted by recency	626	-0.314	1.242	-1.918	1.145	-0.444	0.217	0.141	0.150	0.974	0.573
IV wage bill + experience polynomial	626	-0.376	1.290	-2.001	1.107	-0.583	0.222	0.042	0.163	1.045	0.691
IV wage bill + spell effect	626	-0.284	1.172	-1.792	0.892	-0.507	0.226	0.146	0.183	1.025	0.623
IV wage bill + weighted by recency	626	-0.323	1.234	-1.909	1.066	-0.513	0.210	0.143	0.172	0.996	0.607

Notes: Table reports summary statistics on the ability estimates at the level of the individual worker. Numbers refer to the average ability estimate for a worker across all observed periods. We provide statistics for the full sample and a breakdown by quartiles in estimated worker ability distribution.

Table B2: Correlation among ability estimates at spell level

Correlation coefficient	Add win%	Win% effect	Elo effect	Score dif. + IV	Score dif. + exp.	Score dif. + spell	Score dif. weight	Score dif. + IV + exp.	Score dif. + IV + spell	Score dif. + IV weight	Av. win %	Av. goals scored
Score dif. eff.	0.621	0.900	0.657	0.968	0.961	0.901	0.979	0.931	0.850	0.954	0.592	0.492
Add win%		0.697	0.639	0.613	0.647	0.555	0.607	0.640	0.543	0.601	0.844	0.644
Win% eff.			0.811	0.859	0.869	0.811	0.880	0.830	0.756	0.845	0.669	0.483
Elo eff.				0.623	0.669	0.605	0.658	0.639	0.565	0.626	0.596	0.417

Notes: Table reports the correlation coefficient between different estimation methods for worker 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 worker-firm spell dummies, which are decomposed in a 2nd stage. We do not separately report significance, as this is always found to be below $p<0.01$.

Table B3: Mediocre rehires in alternative estimation models and novice comparisons

Ability threshold	Estimation method	Subsample	Novice comparison	# Spells	# Spells < threshold	% Spells < threshold	# Rehired workers	# workers < threshold	% workers < threshold
Mediocre	Win % eff.	Full	10 years	755	422	55.9%	322	219	68.0%
	Elo eff.	Full	10 years	755	657	87.0%	322	288	89.4%
	Score eff. + Exp.	Full	10 years	755	333	44.1%	322	204	63.4%
	Score eff. + spell	Full	10 years	755	183	24.2%	322	125	38.8%
	Score eff. weight	Full	10 years	755	343	45.4%	322	203	63.0%
	Score eff.	20+ obs.	10 years	679	297	43.7%	275	163	59.3%
	Score eff.	Full	5 years	755	289	38.3%	322	182	56.5%
Sub-standard	Win % eff.	Full	10 years	755	173	22.9%	322	124	38.5%
	Elo eff.	Full	10 years	755	293	38.8%	322	173	53.7%
	Score eff. + Exp.	Full	10 years	755	159	21.1%	322	118	36.6%
	Score eff. + spell	Full	10 years	755	161	21.3%	322	117	36.3%
	Score eff. weight	Full	10 years	755	171	22.6%	322	125	38.8%
	Score eff.	20+ obs.	10 years	679	197	29.0%	275	121	44.0%
	Score eff.	Full	5 years	755	191	25.3%	322	123	38.2%

Notes: See table 4 notes.

Table D1: Linear regression results for probability of employment spell termination in month of analysis

Dependent variable:	Spell Termination		
Ability estimate	-0.026***	-0.022***	-0.025***
	(0.005)	(0.003)	(0.005)
Dif. ability estimate prev. month		-0.033***	-0.038***
		(0.010)	(0.008)
Log tenure			0.003**
			(0.001)
Log age			0.067***
			(0.017)
Log experience			-0.007***
			(0.002)
Constant	0.065***	0.061***	-0.564***
	(0.003)	(0.002)	(0.155)
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

Notes: Table reports linear probability model estimates where the dependent variable equals 1 if the worker leaves the firm 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, *** denotes significance at 1% level, ** at 5% level, * at 10% level.

Table D2: Regression results for career progression after current employment spell

Dependent Variable:	Ever rehired indicator			Career progress		
Ability est. at spell end	0.187***	0.112***	0.110***	0.394***	0.235***	0.235***
	(0.025)	(0.027)	(0.030)	(0.051)	(0.069)	(0.082)
Average win%		-0.095	-0.125		-0.332	-0.360
		(0.201)	(0.181)		(0.523)	(0.466)
Average goals pro		0.249***	0.260***		0.647***	0.681***
		(0.074)	(0.077)		(0.231)	(0.196)
Log games		0.059***	0.061***		0.143***	0.140***
		(0.013)	(0.010)		(0.032)	(0.034)
Log age		-0.583***	-0.591***		-1.447***	-1.598***
		(0.108)	(0.097)		(0.261)	(0.306)
Foreigner		0.225***	0.118**		0.394***	0.168
		(0.052)	(0.051)		(0.143)	(0.149)
Log month number at spell end		0.065*	0.051		0.177*	0.138
		(0.038)	(0.041)		(0.110)	(0.110)
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	0.128	0.134	0.015	0.044	0.047
Estimation method:	Linear Probability Model			Ordered Probit Model		

Notes: The left panel displays regression results for a linear probability model where the dependent variable is an indicator, which equals 1 if a worker is rehired as a manager anywhere in England or abroad after the current employment spell ends, and 0 otherwise. The right panel shows an ordered probit model, where the dependent variable categories are defined as follows, 0 if the worker is never rehired; 1 if the worker is rehired in a lower division as where current spell started; 2 if the worker is rehired in same division as where current spell started, and 3 if the worker 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, *** denotes significance at 1% level, ** at 5% level, * at 10% level.