Choosing the Pond: On-the-Job Experience and Long-Run Career Outcomes

Jie Gong Ang Sun Zhichao Wei*

February 2016

Abstract

This paper investigates the value of on-the-job experience for workers' long-run career outcomes. We exploit the effects of team relegation in professional soccer by contrasting players on teams just below and just above the cutoff point of relegation. We find that players on teams relegated to a lower division have more match appearances in the short run and play in better leagues and earn higher wages in the long run. This gain is concentrated among individuals who are young, or less experienced at the time of relegation. Because other consequences of relegation would tend to work against individuals' long-run success, the positive net gain is most likely to be the result of greater on-the-job experience. The findings have implications for firms on job assignment and for workers on managing their careers.

Keywords: on-the-job experience, job assignment, long-run career outcomes

*Gong: National University of Singapore (email: gong@nus.edu.sg); Sun: Central University of Finance and Economics (email: ang.sun@gmail.com); Wei: China Asset Management (email: zhichaowei@gmail.com). We are grateful to Kenneth Chay, Qiang Fu, Vernon Henderson, Jin Li, Jessica Pan, Ivan Png, Aloysius Siow, Tuan-Hwee Sng and seminar participants at Brown University, Cheng Kong Graduate School of Business, National University of Singapore and Peking University for helpful conversations and suggestions. We thank Shuhe Huang and Jiaxin Ong for research assistance.

1 Introduction

Work experience matters for workers' careers because true skill development happens on the job. A consultant clearly learns more on projects that entail interaction with clients than "on the bench" behind the scenes; a medical intern becomes more proficient by dealing with various complications than by observing cases; and an athlete gains more skills from playing in matches than as a minor substitute. Opportunities to gain such front-line experience are important for human capital accumulation, but are often limited within a firm or organization. Thus, when making exclusive choices of which jobs to take, schools to attend, or teams to join, individuals are concerned with the trade-off between being a big fish in a small pond and a small fish in a big pond. The optimal decision hinges on how much front-line experience matters for their careers in the long run.

To understand the value of on-the-job experience, we examine how it shapes individuals' career development in the long run and whether the effect translates into economically meaningful improvements in labor market outcomes, such as better jobs and higher wages. In an ideal experiment, we would have two subjects with homogeneous abilities and job offers, randomly assign them to workplaces where they would be subject to different levels of on-the-job experience, and analyze their job status and earnings in subsequent years. Realistically, however, this type of research depends on observational data, which suffer from issues of endogeneity. For instance, more capable or motivated individuals are more likely to be assigned to important projects and roles, and thus acquire more valuable experience and enjoy better career outcomes. It is also rare to have a quasi-experiment for estimating the causal effects of experience, which is highly dependent on the joint and nonrandom decisions of the employer and the employee.

The labor market for professional footballers provides a real-world laboratory to examine the value of experience. It is challenging, in most occupations, to observe how much time workers spend performing productive tasks and separate it from general duration of the job. For footballers, the productive task is to play matches and therefore experience can be well

captured by the number of matches played. When a team is relegated to a lower division, it grants more opportunities and playing time to incumbent players, because top talents are no longer affordable. More crucially, relegation is close to a random assignment for teams on the margin, i.e., teams that are just above or just below the cutoff for relegation.

In England, professional football (soccer) is organized as a hierarchy of leagues, with the Premier League at the top, followed by the Football League. Each year, at the end of the season, the three lowest-performing teams in each league are demoted to the next lower division. Teams adjacent to the relegation cutoff are very close in performance and compete fiercely to maintain their top-flight positions. In most years, rankings and relegation results are not finalized until all the matches have been played. It is therefore plausible that players whose teams just missed relegation are comparable to players whose teams were just retained. Relegated teams experience a sharp decline in revenue, reduce their recruiting budgets, and hire less competent players.

By contrasting players on the marginally relegated (third from the bottom) and marginally non-relegated (fourth from the bottom) teams in each season from 1992 to 2002, we show that, first, players on marginally relegated teams were similar to their marginally non-relegated counterparts in demographics, athletic ability, and job characteristics. Second, during the first year after relegation, the appearance rates of the incumbent players increased by 12 percent, relative to the non-relegated players. Third, five to seven years after the shock, players on relegated teams played in significantly better leagues and earned higher wages.

We further find that the effects of team relegation on a player's career are heterogeneous across age groups and prior stock of experience. Only among younger players (between ages 18 and 24) did we find a long-run advantage over non-relegated players. In contrast, older players from relegated teams failed to catch up. The effect also varies by players' prior experience. If players are divided by the median appearance rate in the relegation season, players with less experience show a greater gain from relegation, in long-run jobs

and earnings, while those with an appearance rate greater than the median do not seem to benefit from relegation.

The results are robust to a battery of robustness checks. In particular, we examine sample attrition and player sorting after team relegation. If team relegation tends to shorten the career length of less skillful players, our estimates of the effects of team relegation could be biased upward. If team relegation induces higher turnover and better employee-employer matches for more competent players, our findings could be interpreted as the result of sorting and matching instead of greater experience. We address these issues by analyzing the length of time a player remains in the sample (as opposed to dropping out of the professional leagues) and on the relegated team (as opposed to transferring to another team). We further examine player composition, and find no evidence of a significant difference between relegated and non-relegated teams in player dropout or transfers.

We argue that our estimate is the net effect of team relegation and a lower bound of the effect of on-the-job experience. In addition to more appearances, relegation leads to several other, mostly professionally costly, changes for players. The sources of bias would tend to work against finding positive gains for relegated players. For instance, teams suffer from financial distress and management instability; players may have access to fewer resources and receive inconsistent training; and rival and peer players are less competent than in higher divisions, suggesting smaller spillover from others. Such setbacks in human capital accumulation may counteract the benefits of more on-field experience. Even so, we find that in spite of any losses, relegation can still pay off in the long run for younger players.

We believe that the most plausible interpretation is that for young and inexperienced players, more match experience increases their human capital so significantly that they are better off playing in a lower division than warming the bench in a higher division. Because of financial constraints, relegated teams are less likely to be able to afford players as competent as those they could have attracted if they had not been relegated; therefore, incumbent players are given more opportunity to play in matches. This advantage steadily increases

players' skills and productivity, and appears to improve their long-term value in the labor market. In contrast, if a player had remained in the higher division, he might have spent more time on the sidelines and less time building his human capital.

Our findings have implications beyond professional sports. The "top tier" versus "low tier" leagues capture different levels of prestige in the broader labor market—e.g., an elite law firm versus a small partnership, or a household name consumer packaged goods firm versus a no-name product manufacturer. The "top tier" league corresponds to more prestigious firms, which typically have higher quality peers and more intense internal competition. As valuable experience is usually limited in a firm—as are slots for the first-team squad in a soccer team—for a given worker, the opportunity to acquire valuable experience is slimmer in an elite firm than in a lower-tier firm.

One lesson we draw from the soccer results is on workers choosing jobs. While the very top workers are expected to work in leading places, those who are one step below face a different challenge: If the opportunity to obtain valuable experience is slim in leading firms, they need to figure out whether to stay on the bench or move to a lower-tier workplace to accumulate human capital. Employees facing such trade-offs find themselves in a situation bearing close resemblance to that of soccer players on the margin of higher- and lower-tier leagues. The lesson from our findings is that workers should not automatically assume that they should take the position with a higher-tier firm. The optimal decision hinges upon, among other factors, the degree to which human capital matters for the individual's career path and the extent to which the opportunity to accumulate human capital varies between higher- and lower-tier firms.

Another lesson can be linked to firms' management of human resources. In optimally assigning tasks, firms may want to take into account workers' potential to learn and the effects on their future performance. While delegating the best assignment to the currently most productive worker might be optimal in a static manner, it does not necessarily maximize the firm's long-term value. And, to the extent that workers value experience, firms may attract,

motivate, and retain talents by offering more opportunities to grow(Ke, Li and Powell 2015).

Our findings mainly build on the literature on the value of experience and human capital, a central question in labor economics and personnel strategy(see, for example, Becker 1964). Mincer (1974) and Murphy and Welch (1990) find a positive effect of experience on wages. Topel (1991) and Neal (1995) document the value, respectively, of firm- and industry-specific human capital. Our findings add new evidence that on-the-job experience has positive effects on workers' long-run career outcomes. We show that it is not just job duration, but what kind of task the worker is assigned to, that is important for career development. The experience variable in our study—appearance rate—directly captures the time and intensity of performing tasks. Moreover, since it is unlikely to have a valid quasi-experiment where experience is exogenously changed, our setting is a unique contribution toward estimating how current experience affects future jobs. Lastly, we examine job assignment as a direct outcome, thereby suggesting a channel through which experience affects wages.

By tracing the effects of a discrete setback on individuals' career development, our study is related to the literature on how initial placement and labor market conditions affect workers' long-run careers. Beaudry and DiNardo (1991) find that a worker's current wage is influenced by the lowest unemployment rate since he or she entered the work force. Kahn (2010) and Oreopoulos, Von Wachter and Heisz (2012) find that graduating during a recession leads to lower initial and lifetime earnings. Oyer (2006) shows that among economists, better initial placement leads to long-term benefits in productivity and placement. While the existing literature emphasizes that workers facing unfavorable conditions may be matched to worse jobs and suffer long-run losses, here we show, in a particular labor market, that a negative shock can lead to more opportunities for human capital accumulation, and thereby benefit workers in the long run. Our results suggest that context matters. Caution should be used in applying the mechanism to other settings that may entail career disadvantages.

Professional soccer is an important industry in and of itself. The labor market for soccer players has been widely analyzed by economists and management scholars. Kleven, Landais

and Saez (2013) use data on the international migration of players in European countries and find strong mobility responses to tax rates. Miklós-Thal and Ullrich (2016) study the incentive effects of being selected to Euro Cup national teams. Ichniowski and Preston (2014) track national soccer team performance and the professional leagues their members play for and find evidence of peer effects. Berlinschi, Schokkaert and Swinnen (2013) find that players' migration to foreign clubs improves their origin countries' international performance. Szymanski (2000) use data from English leagues and find evidence of a competitive market for players and racial discrimination in salary.

The rest of the paper proceeds as follows. Section 2 outlines the hierarchical structure of English soccer leagues and describes the impact of relegation. Section 3 describes the data, descriptive results, and empirical strategy. Section 4 presents and interprets the main empirical results. Section 5 reports robustness checks and discusses alternative interpretations. Section 6 concludes the paper.

2 English Professional Soccer

In England, professional soccer clubs compete in a pyramid of leagues, with the Premier League at the top and the Football League—which comprises three divisions—below. Each season lasts for approximately 10 months, from August to May. At the end of each season, a fixed number of the top-ranked clubs from a lower division swap positions with the bottom-ranked clubs in the next higher division.[1] As clubs generate income primarily from broadcast revenue and ticket sales and incur costs mainly on wages, after being relegated they suffer lower revenue and typically resort to cutting recruitment and training.[2] For instance, in the 2010–2011 season, the three clubs relegated from the Premier League—Hull City, Middlesbrough, and Burnley—substantially reduced wages by £17 million (45 percent), £4 million

[1] Appendix Figure A1 depicts the hierarchical structure of the top seven divisions, specifying the number of clubs, promotions, and relegations in each division.

[2] Broadcast revenues are shared exclusively among clubs in the Premier League. For the 2010–2011 season, the league earned total revenue of £2.3 billion, while the second division earned £423 million. See Figure A2.1 for historical revenue trends by league.

(13 percent), and £3 million (13 percent), respectively, while all other clubs in the top and second divisions increased players' wages.[3] Inevitably, relegated teams have difficulty hiring top talents. Statistics show that compared with non-relegated teams, the relegated teams hired players from lesser leagues.[4]

Professional soccer players are highly trained athletes who exert effort to excel at the game. They sign contracts with clubs, specifying duration and annual salary.[5] During competitions, a player must constantly process information about his position on the field, coordinate with peers, and react to rivals' moves—activities that require skills and experience and lead to improved performance. These skills are primarily developed from training and playing in practice games and matches.[6]

Relegation is generally believed to be professionally costly for the affected players. As relegated clubs play in lower-level divisions and cut salaries, their players are likely to experience wage cuts and a decline in the quality of training. Tight budgets restrain these clubs from recruiting top players, and hence the retained players enjoy smaller spillover of skills from their peers. Compared with the top division, they also play against less competent rivals, which means smaller spillover from players in other clubs. The players also receive less attention from the media and scouts.

However, relegation could bring an important benefit—more match appearances—for the retained players. When clubs move to lower divisions, they can no longer afford top talents;

[3]Source: Annual Review of Football Finance, Deloitte 2012.

[4]We take players hired within one year after relegation and compare their previous league levels. On average, the relegated teams hired new players from league level 2.77 while the non-relegated teams recruited from league level 1.67 (smaller numbers represent better leagues). A t-test of the means rejects the null hypothesis that the recruits' prior league levels are equal for the relegated and non-relegated teams (p-value=0.05, two-sided t-test).

[5]Contracts are typically multi-year. If a player wants to move to a new club, the two clubs negotiate a transfer fee that the destination club pays to the incumbent club.

[6]As a comparison, consider production-line jobs that are relatively routine and could be easily learned. One such example is a production job at Ford. Raff and Summers (1987) provide the following quote from Meyer (1981, p. 41): "Division of labor has been carried on to such a point that an overwhelming majority of the jobs consists of a very few simple operations. In most cases a complete mastery of the movements does not take more than from five to ten minutes. All the training that a man receives in connection with his job consists of one or two demonstrations by the foreman or the workman who has been doing that job. After these demonstrations he is considered a fully qualified 'production man.' All that he has to do now is to automatize these few operations now so that speed may rapidly be increased."

non-star players, therefore, may move up in internal rankings and are more likely to play in league competitions. For a professional athlete whose career objective is to win league and international tournaments, actual competitions provide the best opportunities to accumulate on-the-job experience and human capital. Understanding this, top-flight clubs often send young players to lesser clubs to gain first-team experience. Sophisticated agents also suggest that young players choose middling clubs rather than being a minor substitute in elite clubs.[7] Voluntarily "moving down" to a lesser league is an endogenous decision made by players and agents, and therefore cannot be used to estimate the causal effect of being assigned to a lesser league. Our identification relies on marginal relegation, that is, players in teams just above and below the cutoff point, which is an unintended shock to the concerned workers.

3 Data and Empirical Strategy

3.1 Panel Data for Player Appearances and Assignments

We compiled a panel dataset of 664 players from multiple sources.[8] We first collected the marginally relegated (third from the last) and marginally non-relegated (fourth from the last) clubs in the Premier League for each season from 1992 to 2002.[9] Then, for every player on the clubs' squads at the season of relegation, we collected their demographics and complete career history. Player information includes name, date of birth, club affiliation and appearance rate for each season, and wages. About 73 percent of the players were from England. The average career length was 16.78 years, during which the player transferred across clubs 5.48 times.

[7]For example, Brazilian star Bernard chose Porto over the elite Arsenal. His agent commented that "this would be the best solution for him. He would prefer to join a middling, but good, club in Europe rather than risk joining an elite club and not getting the chance to play." Source: ESPN, July 29, 2013.

[8]The main online sources are the websites playerhistory.com and premierleague.com.

[9]We focus on relegation from the Premier League, instead of promotion or relegation between other lower levels, because the data are most complete for players and teams at the Premier League. The sample starts with the inaugural season 1992–1993 and stops at season 2002–2003, so that the last group of relegated players still have 10 years of data after, allowing us to identify any long-run effects.

A key outcome variable is a player's job assignment, namely, the level of the league for which he plays. About 88 percent of player-season observations are associated with English leagues and therefore the job assignments are naturally ranked by the order of the leagues: Premier League = 1, Football League Championship = 2, Football League One = 3, and so on. The remaining 12 percent of the observations involve non-English leagues. To make non-English leagues comparable to English leagues, we reweight foreign leagues' levels using the Union of European Football Associations' country coefficients, which rank the relative strength of national leagues each year.[10]

Player-level wages are generally confidential and publicly unavailable, except for a few star players. The individuals in our sample were typically non-top players who were playing for clubs on the margin of relegation, and therefore their wages were not revealed publicly. We use a novel data set from *Football Manager*, a simulation game based on real player information. The game collects information from scouts and journalists and sells the data set to prestige clubs.[11] The wage data we extracted have two limitations. First, there could be measurement error. This would only attenuate our estimates if it is not systematically correlated with relegation. Second, the missing data on wages are not negligible. In the robustness checks, we address this concern and show that our main results are not driven by the attrition of wages.

3.2 Random Assignment of the League Level

Relegation is close to a random assignment of league level in two ways. First, the marginally relegated and marginally non-relegated teams are close in performance and in player composition. Figure 1 shows the teams' points at the completion of each season (1980–2010) for the champion, marginally non-relegated (fourth from the last), marginally relegated (third from

[10]For example, for the 2012–2013 season, Spain topped the coefficient ranking with a coefficient of 17.714, while England's coefficient was 16.428. Therefore, relative to England's Premier League (level = 1), Spain's top division La Liga in 2012–2013 was at level $1 \cdot (16.428/17.714) = 0.927$.

[11]Clubs use the *Football Manager* data set to research players. For instance, in 2008 the premiership team Everton signed a contract with Sports Interactive, the studio behind the game, to use the database to scout players and staff.

the last), and bottom-ranked teams. In all years, points are very close between marginally relegated and marginally non-relegated teams, in contrast to the significant disparity between the top- and bottom-ranked teams.

[Figure 1 here]

More importantly, players from the marginally non-relegated teams make a comparable control group for those from the marginally relegated teams. We compare predetermined player characteristics between the two groups before the relegation season. In Table 1, columns (1) and (2) show the means of predetermined player characteristics for the relegated and non-relegated teams, respectively. Column (3) reports the differences. Relegated players are similar to non-relegated players in terms of demographics (height, weight, nationality, and age) and professional characteristics (years in the sport, loan or swap history, wage, and transfer fee). For each of the variables, a t-test cannot reject the null hypothesis that the means are the same between the relegated and non-relegated groups. It therefore suggests the two groups' resemblance, to the extent that the means of observable characteristics are balanced.

[Table 1 here]

Second, it is unlikely that teams manipulate the final ranking or choose strategies based on different expected likelihoods of relegation. For teams on the margin of relegation, the competition for survival is so fierce that relegation is usually not determined until all the matches have been played. In 10 of the 11 seasons from 1992 to 2002, for instance, relegation was only determined in the final week, during which the survivors (marginally non-relegated teams) won by small margins.[12]

[12]The only exception was season 2000–2001, in which relegation was confirmed in the penultimate game.

3.3 Mean Comparison and Graphical Evidence

In this subsection, we compare the means of appearance rate and league level for players from marginally relegated versus marginally non-relegated teams. A player's appearance rate is the ratio of the number of matches in which he appeared on the start-up squad to the total number of matches his team played in the season.[13] The league level is defined such that the Premier League is level 1, the Champion League is level 2, etc., and higher numbers indicate lesser leagues.

Figure 2 presents the means of appearance against year relative to relegation, running from the fourth year before to the tenth year after relegation. Figure 3 presents the means of players' league levels. We inverted the vertical scale so that smaller numbers represent higher-ranked (better) divisions.

[Figure 2 here]

[Figure 3 here]

Figures 2 and 3 suggest two empirical regularities. First, the pre-trends of appearance rates and league levels are very similar for marginally relegated and marginally non-relegated players. Second, as seen in Figure 2, relegated players appeared substantially more often than non-relegated players immediately after relegation. In comparison, Figure 3 shows that in the first three years after relegation, relegated players belonged to worse leagues than their non-relegated counterparts, but this gap narrowed over time and the trend was eventually reversed. In the fourth year after relegation, relegated players began playing in better leagues than non-relegated players.

These findings are consistent with our premise that relegated players had more opportunities to play on the field, which helped them develop skills and resulted in better job assignments in the long run. Intuitively, the effect of human capital accumulation should

[13] When a player played in multiple divisions during one season because of loan, swap, or transfers, we assume that he spent equal time in each division and calculate the average appearance rate.

vary by a player's age; any on-field experience that affects skills and future assignments would have a larger effect on younger players. Figures 4.A and 4.B show that better future assignments for relegated players are only observed among those ages 24 years or younger; older players gained no clear short- or long-term advantage by being relegated.

[Figure 4.A and 4.B here]

To see the magnitude and significance levels of the short-run versus long-run effects, we summarize the means of player appearance and job outcomes at different periods around relegation. As shown in Table 2, columns (1) to (3), before relegation, the means of appearance rate, league level, and wage are similar between the two groups of players. Immediately after relegation, the relegated players play in worse leagues and appear in more matches than the non-relegated players do (columns (4) and (5)). The difference in appearance rate is about 0.07 (14 percent increase compared with the non-relegated players) and is significant at the 1 percent level. During the five to seven years after relegation, the relegated players play in better leagues and earn higher wages (measured as weekly income in pounds) than non-relegated players (columns (7) and (8)), and the difference is significant at the 5 percent level. This simple mean comparison suggests that relegated players appear in more matches during their stay in the lower division, and play in better leagues in about five years.

[Table 2 here]

3.4 A Differences-in-Differences Approach

For the econometric analysis, we use a differences-in-differences (DID) approach to exploit the variation in time elapsed since the relegation season and cross-sectional variation in whether a player was exposed to team relegation. Our main analysis focuses on the interaction of time and cross-sectional variations. The DID estimation will remove any possible time-invariant

heterogeneity between the two groups of players. In addition, it allows us to examine short- and long-term effects in the same regression. The econometric specification is as follows:

$$Y_{it} = \alpha + \sum_{\tau=1}^{10} Relegated_i \cdot Period_\tau \cdot \beta_\tau + \sum_{\tau=1}^{10} Period_\tau \cdot \gamma_\tau + \theta_i + \delta_t + \varepsilon_{it} \tag{1}$$

where i and t index players and time periods, respectively. Y_{it} represents three main outcomes for player i at period t: his annual appearance rate, affiliated league level, and weekly wage. $Relegated_i$ is a dummy variable that equals 1 when player i is marginally relegated. $Period_\tau$ are dummy variables for time periods relative to the season of relegation. We focus on five seasons before and ten seasons after relegation. We group the seasons into four time windows and report average effects for each: one year after, two to four years after, five to seven years after, and eight to ten years after relegation. The reference group is comprised of players' job history from the fourth year before to the year of relegation. To filter the variations in productivity driven by the life cycle, we include a third-order polynomial function of player age. θ_i and δ_t are the player and calendar year fixed effects, respectively. Lastly, ε_{it} is the error term and is clustered at the player level in all regressions.

The coefficient of interest is β_τ, the mean shift in appearance rate and assignment after relegation among players in the marginally relegated teams, relative to the marginally non-relegated players. β_τ can be interpreted as the causal effect of relegation, under the assumption that the differences between relegated and non-relegated players are time-invariant. The identification assumption is supported by the pre-trends ploted in Figures 2 and 3, in that for every season preceding relegation, the two groups of players had very similar appearance rates and played in similarly ranked leagues. We further conduct a robustness check of parallel pre-trends. We regress outcomes—appearance rate, league level and wage—in pre-relegation seasons on a linear time trend, the relegation dummy, and their interactions. The coefficients of the interaction terms can be interpreted as differences in pre-trends. We find that none of the coefficients is statistically significant, suggesting that one cannot reject the null hypothesis that the two groups of players had parallel pre-trends.

3.5 Identification Particulars

We discuss three particulars of our DID estimations: sample attrition, rotating panel data and recurring players.

Sample attrition. If relegation affects players' career length and, therefore, the length of time they remain in the sample, our DID estimation will be biased by sample attrition. For example, if low-skilled players on relegated teams are more likely to drop out because of dimmed career prospects, we will overestimate the effect of relegation. We formally test whether relegation affects the length of time a player remains in the sample, and further compare the means of dropouts' characteristics between the relegated and non-relegated groups. We find that there is no significant difference in the duration of remaining in the sample and that the sample means of the dropouts' characteristics are similar across the two groups of teams, which suggests that sample attrition is unlikely to drive our results. We report the detailed results in section 4.3.

Rotating panel data. We construct the panel data by identifying players who were once on the margin of relegation and tracking their career histories. Because they started and ended their professional careers at various times, the composition of players in our data changes over time. For example, a player who was 30 years old when relegated might remain in the sample for five years and have a wealth of pre-relegation data. In contrast, a player who was 20 when relegated might remain for 10 years, but virtually have no data before relegation. These differences will not affect the DID identification, however, under the assumption that relegation does not affect sample attrition. That is, if the rotating patterns are similar between the relegated and non-relegated groups, heterogeneity caused by attrition will be canceled out by DID estimation.

Recurring players. In our sample, about 10 percent of the players experienced the event more than once—i.e., they had been on a third-last or fourth-last team multiple times during their careers. When organizing the panel data, we regard each instance as a different individual. That is, the second time a player is on the margin of relegation, we assign him

a new ID and treat him as a different player. We use the number of times a player has been on the margin as the weight in our regression analysis. Note that given the random nature of relegation for teams on the margin, a retained player may eventually be marginally relegated (either for the same team or a different team) and vice versa. This means that the control group could be contaminated, and therefore the effects of relegation could be underestimated.[14]

4 Results: Effects of Relegation

4.1 Effects of Relegation on Experience and Job Outcomes

Table 3, columns (1) and (2), reports the effect of relegation on players' annual appearance rate. Column (1) presents the DID estimates, controlling for player age and season fixed effects (e.g., the 1992–1993 season), and column (2) further controls for player fixed effects. Both estimates are consistent with Figure 2, showing that the annual appearance rate increased immediately after relegation. The estimates are statistically significant at the 5 percent level. Using the most conservative estimates, on average, players on relegated teams appeared 5.9 percent more often, or in 2.24 more matches, during the season immediately after relegation. Evaluating at the sample mean (51 percent), relegation increases players' match appearances by about 12 percent.

[Table 3 here]

Table 3, columns (3) and (4), reports the effects on the players' league level. A smaller value for league level means a better division; hence, negative coefficients indicate playing in better leagues. The initial effect of relegation on league level is positive (lower level), but

[14] A given team may have been ended at the margin of relegation more than once too. This is not a problem for our identification because we track players instead of teams. By the next time a team ended at third or fourth from last, many players from its original squad had already left. Since 90 percent of the players have been on the margin only once, the estimate is the effect of relegation, rather than the effect of being relegated earlier versus later.

the effect becomes negative (higher level) around the second to fourth years after relegation. During the fifth to seventh year, the effect becomes very significant. These findings suggest that in the short run, the relegated players are stuck in a lower division, but in the long run, they play in better leagues than the non-relegated players.

To round out our analysis of the impact of relegation on players' careers, we examine player wage dynamics. As reported in Table 3, column (5), starting around the fifth year after relegation, relegated players on average earned more than non-relegated players. The coefficients are significant at the 10 percent level. Adding player fixed effects (column (6)) lowers the identification power by dropping variations between players. Overall, the wage patterns are in favor of the value of on-the-job experience.

Lastly, we address the concern about missing wage data. If the incidence of missing value is not random—for example, collecting wages from lesser leagues could be so costly that journalists and scouts do so only for the top players in the relegated teams, while all players in the Premier League have complete records of their wages—then the wage effect we find could be driven by sample selection. We formally test the effect of relegation on the incidence of missing values on wages. We use a dummy variable indicating whether player i's wages in season t are missing as the dependent variable, and fit in regression (1). As reported in Appendix Table A5, there is no significant difference in the likelihood of missing wage data between the relegated and non-relegated players.

4.2 Evidence for Human Capital Accumulation: Heterogeneous Effects across Player Age and Prior Experience

Our main premise is that relegation increases players' chances to play on-field and therefore helps them accumulate human capital. Intuitively, the effect of experience on long-run career should vary by player age. Age is critical to a soccer player's productivity due to the raw physical demands of the sport. As a player ages, his physical conditions declines and injuries accumulate, limiting the extent to which he can capitalize on the skills and techniques he

has learned. There is also evidence that human capital investment has diminishing returns (Mincer 1974; Murphy and Welch 1990), so investments that affect skills should have a larger return for the younger players.[15]

A testable implication is that younger players should benefit more from on-field experience than older players. We divide players into two subsamples, those between ages 18 and 24 years, and those between ages 25 and 30 years at the time of relegation. Table 4, columns (1) and (2), shows the effects on league level for younger and older players separately. We observe that, from the fifth to seventh year after relegation, the younger relegated players were in significantly better leagues than the non-relegated players. In contrast, the effects on older relegated players are much smaller and not significant.[16]

A similar argument can be made about prior experience. As investment in human capital has diminishing returns, playing matches should have a stronger effect on players who have had less playing opportunity than on those who have already shown up in many matches. We divide players by the appearance rate at the relegation season. The sample median is around 0.52. Table 4, columns (3) and (4), shows the impact of relegation on players above and below the median appearance rate, respectively. The contrast is very clear: the relatively inexperienced players—individuals who previously appeared in fewer games—show a significant gain from relegation; they play in leagues that are more than one division higher than the non-relegated players. Players with more prior experience, however, do not seem to benefit from relegation; the coefficients are close to zero and not significant.

[Table 4 here]

Table 5 reports the heterogeneity of the effects on wages. As shown in columns (1) and

[15]Another reason is that younger players are more motivated to accumulate human capital, because they can reap the benefits over a longer horizon (Becker 1967; Ben-Porath 1967). Therefore, a younger player may exert more effort for the same amount of playing time than an older player.

[16]We test the effect of relegation on appearance rate for young and old players separately and find that they both receive more chances to play after relegation. Hence, the heterogeneous effects are not caused by different treatment intensity.

(2), players who were younger (between 18 and 24) when relegated earned higher wages in the long run than non-relegated players. However, we do not find such wage gains for relegated players who were age 25 or older at the relegation season. The contrast is also found across players with different prior appearance rates (columns (3) and (4)). Higher earnings after relegation are only found among individuals whose prior appearance rate was below the median.

[Table 5 here]

We test whether the estimates are statistically different between the subgroups. As shown in Appendix Table A7, the differences in the effects on league levels (column 3) and wages (column 4) between more experienced and less experienced players are both significant at 10 percent level. For the comparison between young and old players, the difference in the wage effect is significant at 10 percent level (column 2). While the difference in the effect on league level is not statistically significant (column 1), the sign is consistent with our conjecture that young players benefit more from relegation. In summary, the differential patterns between age and experience subgroups are in favor of human capital accumulation from more playing time.

To the extent that players benefit—especially the younger and inexperienced ones—from relegation, one might be concerned that they might withhold effort to increase the chances that their team is relegated. The benefit from relegation notwithstanding, we believe that this type of moral hazard problem is unlikely to exist in practice. Soccer is fundamentally a team sport, and it is less clear how a single player's performance determines the final outcome. More importantly, players' actions are clearly observed. Egregious actions of low effort are easily detected and punished, and low-level performance sends bad signals to the market, in the sense that the market updates downward the players ability. Therefore, the benefit of shirking is small for the players and the cost is remarkably large. This suggests that the moral hazard problem of the players is of second importance in these games.

4.3 Effects of Relegation on Player Dropout and Transfer

In this subsection, we analyze the effects of relegation on players' decisions to drop out of professional leagues and transfer to another club. Sorting following negative career shocks is an interesting outcome in its own right, but more importantly, it may affect our estimates and the interpretation of the results.

The abovementioned estimates are the average effects of relegation on all players in the team, including players who were retained and those who left the team months or years after relegation. Appendix Figure A3 plots the percentage of retained players by time and relegation outcome. Overall, among players on the squad at the relegation season (year 0), 70 percent remained on the team one year after relegation (year 1); the share drops to 50 percent within another year (year 2). The dynamics are similar across marginally relegated and marginally non-relegated teams. This means that around 70 percent of the players are retained in the year immediately after relegation and experience the adversities of the lower division. The other 30 percent of the players either dropped out of the professional leagues (and therefore our sample) or moved to different clubs. While dropping out of our sample may bias our estimates, transferring to other clubs may affect the interpretation of the results.

First, our estimates could be biased because of sample attrition and the direction of the bias can go either way. It is possible that once relegated, lower-skilled players are more likely to drop out because of dimmed career prospects. In that case, we overestimate the effects of relegation. In contrast, if more skillful players are more likely to drop out as outside options (e.g., coaching or managing other clubs) yield higher returns than resuming their careers as footballers, the effects of relegation could be underestimated.

We directly test the effect of relegation on the length of time a player remains in the professional league, and therefore in our sample. Appendix Table A1, column (1), reports the results from Poisson estimation. The coefficient is small and not significant, suggesting that attrition from the sample is not affected by relegation. We further examine the characteristics

of players who dropped out, as different dropout compositions could also bias our estimates. For example, even if, on average, the two types of teams had similar dropout rates, it is possible that they lose different types of players—e.g., relegated teams maintained less capable players than the non-relegated teams, in which case the estimates would be upward-biased. Appendix Table A2, panel A, presents the mean comparison of the observable characteristics for players who dropped out within five years after relegation. The results suggest that the characteristics of the dropouts from relegated and non-relegated teams are balanced, in that we cannot reject the null hypothesis that the mean is equal between the two groups of players for most of the observable variables.

Second, it is possible that relegation triggers higher turnover and facilitates player-club matches, and therefore the relegated players have better jobs and higher earnings. For instance, a relegated club might sell its most expensive players for cash to cope with financial distress; in that case, the observed better assignments were likely driven by the top talents transferring to elite clubs. We address this potential problem by studying the effect of relegation on transfers. Appendix Table A1, column (2), shows the Poisson estimate of the effects of relegation on the length of time a player remained on the team until he transferred out. The coefficient is small and not significant. This implies that relegated players were neither more nor less likely to transfer to other clubs. Appendix Table A2, Panel B presents the mean comparison of observable characteristics for players who transferred to other clubs within one year after relegation. Most of the characteristics are well balanced between the relegated and non-relegated players.[17] We cannot reject the null hypothesis that the two groups are similar regarding the characteristics of players who transferred out. It is therefore unlikely that our findings are driven by player sorting after relegation.

We conduct an additional robustness check by dropping players that moved to another team directly after the relegation season. Appendix Table A3 reports the results. Overall, the estimates are similar to those from the baseline specifications. Some estimates have

[17]The only exception is player height. Despite its statistical significance, the mean difference is only 0.01m, very small compared with the group means (1.80m).

slightly larger standard errors, partly because of the smaller sample size. The main pattern remains robust: Compared to non-relegated players, relegated players who remain on the team appear in more matches and, starting from the fifth year after relegation, play in higher-ranked leagues.

5 Robustness Checks

5.1 Design Validity: Exclusion of Predictable Relegation

Our research design relies on the assumption that players from marginally relegated and marginally non-relegated teams are comparable, except for the assignment of league level. One threat to our research design is that the relegated teams had different prospects and strategies from the non-relegated teams. For instance, teams may be promoted to the top flight by luck or some other one-time event. Realizing that their chance of survival is slim, these teams might spend less on players during their brief stay in the top flight than they spent during their time in the lesser league (Noll 2002).

To address this issue, we identify such predictable relegations from teams' historic performance in the leagues. We drop four teams that were promoted to the Premier League, stayed for only one season, and were immediately demoted. Appendix Table A4 reports the estimates for the remaining sample. The results are very similar to the benchmark estimates in Table 3, suggesting that our findings are not driven by teams' strategic behavior.

5.2 Alternative Explanations

Our main results show that relegation leads to an immediate increase in appearance and long-run advantage in jobs and wages for the players. We interpret the link between appearance rate and better career outcomes as evidence that experience improves human capital. Here we discuss three competing explanations: recruiters' bias, player visibility, and incentive for promotion.

The first alternative explanation is that players are judged by their quality relative to those on the pitch with them. After relegation, a high-ability player who finds himself surrounded by poor players looks "too good" compared with others in the league, and therefore he looks attractive to those recruiting in higher leagues. Recruiters might suffer from behavioral bias, overvalue match performance and fail to adjust for weaker competition.

We first collect data on the relegated teams' performance in the second division (Football League Championship). It appears that the relegated teams did not always look too good relative to the other teams in the second division. Their final points are between -0.8 to 2.33 standard deviations from the mean. In several seasons, the relegated team was ranked very low. Then, we conduct a formal test by separating players who may look too good from those who may not look too good. We use the Football Manager's technique index and divide the sample by the median. If recruiters' bias drives our results, then the benefits from relegation would be concentrated on those who look too good, i.e., more skillful, players. As shown in Appendix Table A6, the long-run benefits in league level is also found among low-skilled players—those who do not look too good. Lastly, we argue that recruiters bias may benefit players in the short run, but the effect is unlikely to persist in the long run. Unless they improve their skills, once the relegated players move to higher-ranked leagues, they no longer look too good relative to those on the pitch with them. The market should update the players' ability downward and the initial bias will be corrected. This behavioral hypothesis therefore cannot fully explain the benefits starting from the five to seven years after being demoted.

The second alternative explanation is that match appearances expose the player to media, professional agents, and other clubs—and, potentially, more job opportunities. Visibility is particularly important for showcasing a player's talents if potential employers are risk-averse. By appearing in more matches, players reveal more information about their skills and ability, and the risk of a bad match between player and club is smaller.

We cannot completely separate experience from visibility, as they move together with

match appearance. While experience may improve player skill in a persistent manner, visibility mainly facilitates employer learning when the market has limited information about the player. One implication is that the effect of visibility should diminish as the employer learns more about players' skills from daily training and performance. If relegated players moved to better clubs because of sheer visibility, we should see the largest effect in the first few seasons after relegation, and a smaller or no effect in the long run. However, our findings show the opposite trend: more appearances in the first year did not lead to better jobs until five to seven years later, by which time an average player had transferred four times. The dynamics are more consistent with a persistent increase in player skills than employer learning based on one-season observations.

A third possible competing mechanism is the incentive to obtain promotion to the Premier League. Relegated players may increase their level of effort to obtain promotion back to the top division and higher effort may also improve skills. Anecdotal evidence suggests, however, that the relegated players did not exert great effort during their stay in the Football League. Given that they have played in the most competitive league, we would expect them to lead by a large margin in the lower league. However, among the eleven seasons in our sample period, the relegated teams made into the top three only in five seasons. They were even ranked below the eighth in three seasons. It seems that the team morale was low and the motivation for promotion was not particularly strong.

6 Conclusion

We investigated the returns from experience by estimating the effects of team relegation from the English Premier League on players' future jobs and wages. We contrasted players on marginally relegated and marginally retained teams and found unintended, positive consequences of relegation: players relegated to the lesser division, especially those who are younger or inexperienced at the time of this career shock, do better in the long run than

players remaining in the top division. Since most of the heterogeneity induced by relegation tends to work against the relegated players, we conclude that the improved job assignments and higher wages are the result of the greater experience immediately following relegation.

Our findings suggest that on-the-job experience matters. In particular, performing the most relevant task intensively benefits workers in the long run. This lesson could be applied, in principle, to situations such as when a worker fails to join an elite firm, or a student fails to be admitted to a top university. If in less prestigious organizations, they stand out from their peers and have more opportunities to practice, this is not necessarily a bad thing in the long run.[18] A caveat in inferring broader lessons is that the worker accumulates human capital through performing the same task. Soccer players do the same task—playing matches—in either the top- or second-tier leagues, so by having more playing time, they build skills for that specific task. This is an important feature of the soccer industry. If joining a less prestigious organization instead involves different tasks, our mechanism may not apply and different outcomes may occur.

References

Beaudry, Paul and John DiNardo, "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data," *Journal of Political Economy*, 1991, *99* (4), 665–688.

Becker, Gary S, "Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education," 1964.

_ , *Human Capital and the Personal Distribution of Income* 1967.

Ben-Porath, Yoram, "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 1967, *75* (4), 352–365.

[18]For instance, Conley and Onder (2013) find that in the labor market for young economists, the best graduates at non-top-ranked institutions have better publications than the median graduates from top schools.

Berlinschi, Ruxanda, Jeroen Schokkaert, and Johan Swinnen, "When Drains and Gains Coincide: Migration and International Football Performance," *Labour Economics*, 2013, *21* (C), 1–14.

Conley, John P and Ali Sina Onder, "An Empirical Guide to Hiring Assistant Professors in Economics," 2013. Vanderbilt University Department of Economics Working Papers, 13-00009.

Ichniowski, Casey and Anne Preston, "Do Star Performers Produce More Stars? Peer Effects and Learning in Elite Teams," 2014. NBER Working Paper No. 20478.

Kahn, Lisa, "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labour Economics*, 2010, *17* (2), 303–316.

Ke, Rongzhu, Jin Li, and Mike Powell, "Managing Careers in Organizations," 2015. Working paper.

Kleven, Henrik Jacobsen, Camille Landais, and Emmanuel Saez, "Taxation and International Migration of Superstars: Evidence from the European Football Market," *American Economic Review*, 2013, *103* (5), 1892–1924.

Meyer, Stephen, *The five dollar day: Labor management and social control in the Ford Motor Company, 1908-1921*, State University of New York Press, 1981.

Miklós-Thal, Jeanine and Hannes Ullrich, "Career Prospects and Effort Incentives: Evidence from Professional Soccer," *Management Science*, 2016.

Mincer, Jacob, *Schooling, Experience, and Earnings*, Cambridge, MA: National Bureau of Economic Research, 1974.

Murphy, Kevin M and Finis Welch, "Empirical Age-Earnings Profiles," *Journal of Labor economics*, 1990, pp. 202–229.

Neal, Derek, "Industry-Specific Human Capital: Evidence from Displaced Workers," *Journal of Labor Economics*, 1995, *13* (4), 653–677.

Noll, Roger G, "The Economics of Promotion and Relegation in Sports Leagues The Case of English Football," *Journal of Sports Economics*, 2002, *3* (2), 169–203.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz, "The Short-and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.

Oyer, Paul, "Initial Labor Market Conditions and Long-Term Outcomes for Economists," *Journal of Economic Perspectives*, 2006, *20* (3), 143–160.

Raff, Daniel M.G. and Lawrence H. Summers, "Did Henry Ford Pay Efficiency Wages?," *journal of Labor Economics*, 1987, *5* (4 pt 2), S57–S86.

Szymanski, Stefan, "A Market Test for Discrimination in the English Professional Soccer Leagues," *Journal of Political Economy*, 2000, *108* (3), 590–603.

Topel, Robert, "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority?," *Journal of Political Economy*, 1991, *99* (1), 145–176.

Figure 1. End-of-Season Points of Premier League Clubs 1980-2010

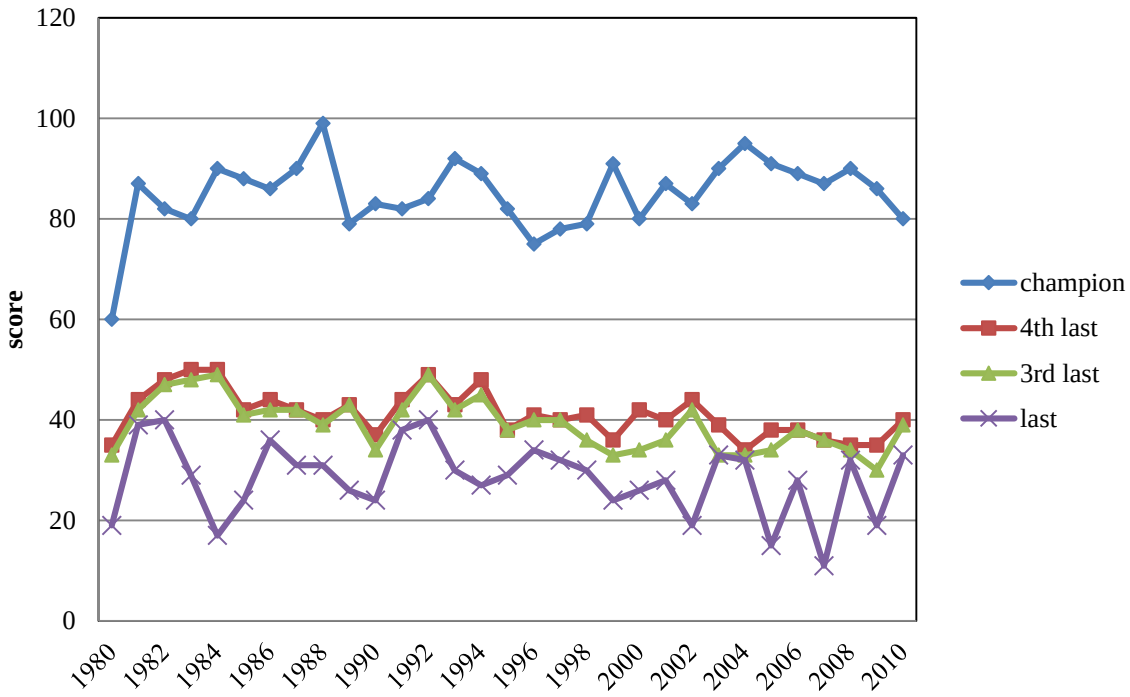

Source: Premier League final score and ranking tables are online at http://www.premierleague.com/en-gb/matchday/league-table.html

Notes: The figure depicts the points of clubs with different standings, namely, the champion, the 4th last (i.e., the club just above the cutoff), the 3rd last (i.e., the club just below the cutoff), and the bottom-ranked at the end of each season between 1980 and 2010. Vertical axis is clubs' end-of-season points, based on their match results during the season. A club earns 3 points for winning a match, 1 point for a tie, and 0 for losing.

Figure 2. Sample Means of Annual Appearance Rate by Relative Years

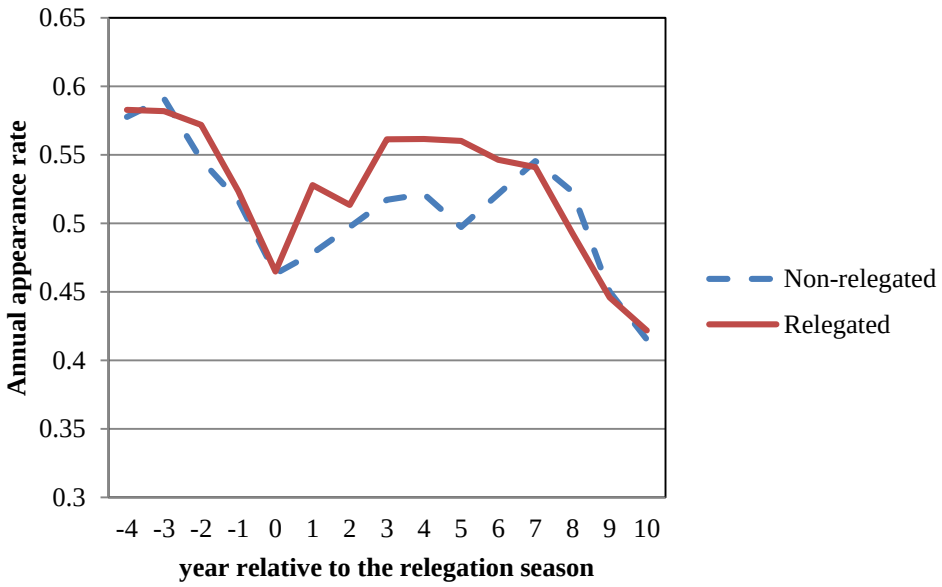

Notes: This figure shows the sample means of annual appearance rate by year relative to the relegation season for relegated and nonrelegated clubs. Horizontal axis is years leading up to relegation (negative) or since relegation (positive). Vertical axis is annual appearance rate, calculated as the ratio of the number of matches in which a player appeared on the start-up squad to the total number of league competitions the club played in a season.

Figure 3. Sample Means of League Levels by Relative Years

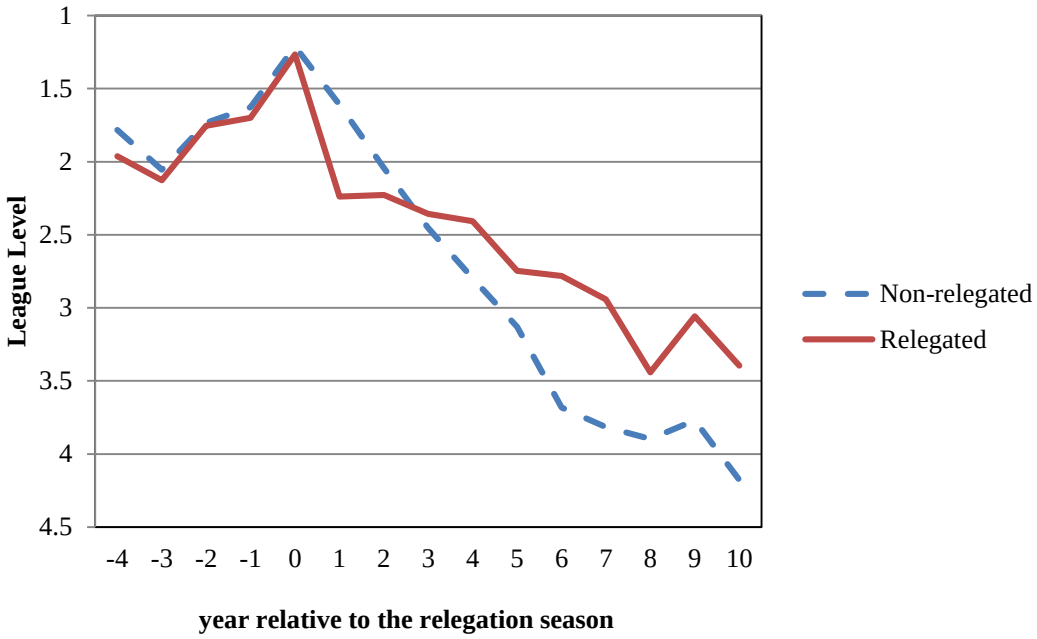

Notes: This figure shows the sample means of league levels by year relative to the relegation season for relegated and nonrelegated clubs, respectively. Horizontal axis is years leading up to relegation (negative) or since relegation (positive). Vertical axis is the level of leagues, with the highest being Premier League = 1 and the lowest Football League 2 = 4. Non-English leagues are reweighed by the ratio of England's UEFA country coefficient to the respective country's country coefficient.

Figure 4A. Heterogeneity across Age: Means of League Level for Players below 24 at Relegation

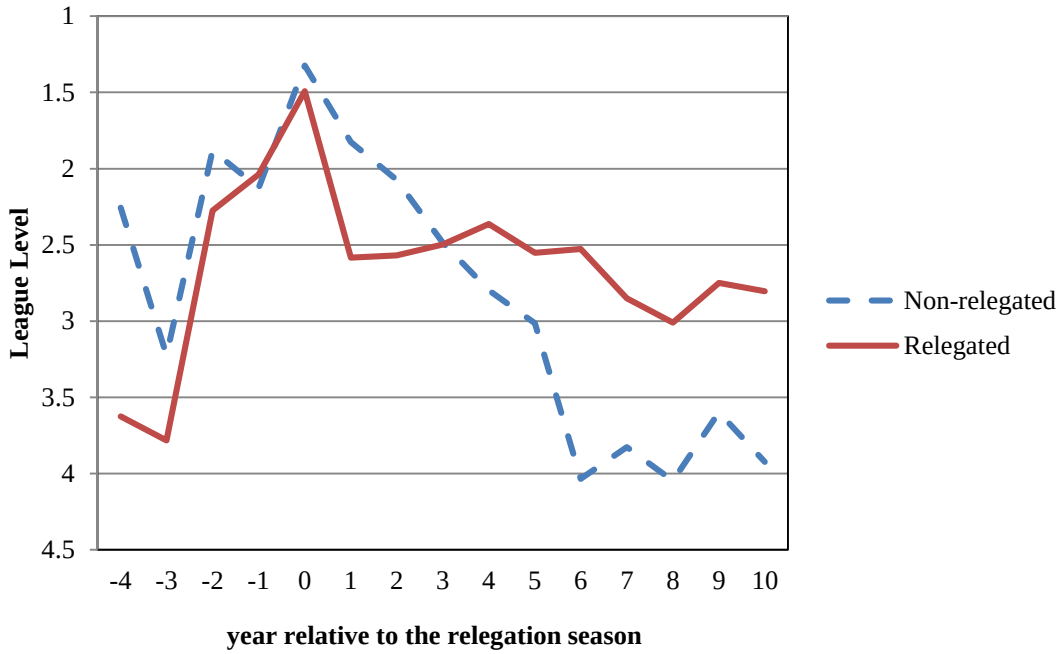

Figure 4B. Heterogeneity across Age: Means of League Level for Players above 24 at Relegation

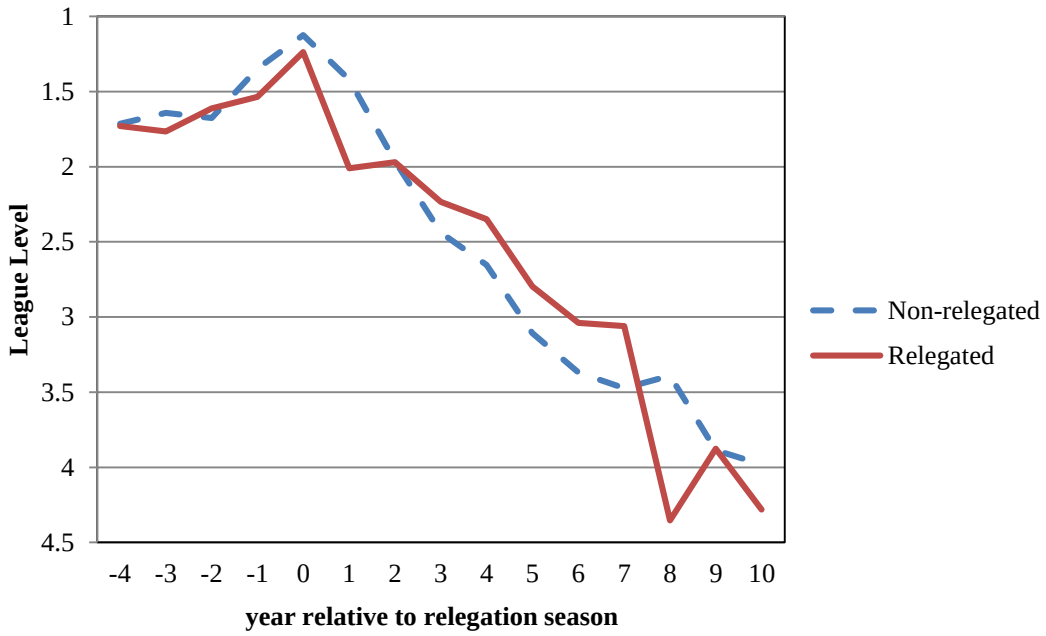

Notes: These graphs depict the means of league level by year relative to relegation for relegated and nonrelegated players, and how this differs by player age. Figure 4A uses the subsample of players younger than or equal to 24 years of age at relegation and Figure 4B uses the sample of players older than 24 at relegation. Horizontal axis is years leading up to relegation (negative) or since relegation (positive). Vertical axis is the level of leagues, with the highest being Premier League=1 and the lowest Football League 2 = 4. Non-English leagues are reweighed by the ratio of England's UEFA country coefficient to the respective country's country coefficient.

Table 1. Predetermined Player Characteristics by Relegation

Predetermined player characteristics	Relegated club players	Nonrelegated club players	Difference (1)-(2)
	(1)	(2)	(3)
Height (m)	1.808	1.801	0.006
	(0.063)	(0.068)	(0.005)
Weight (KG)	76.772	76.341	0.432
	(7.415)	(8.716)	(0.622)
Foreigner (Yes=1, No=0)	0.228	0.228	(0.000)
	(0.420)	(0.420)	(0.031)
Age at relegation season	25.591	25.511	0.080
	(4.993)	(5.254)	(0.378)
Years of experience before relegation	7.860	7.522	0.338
	(5.006)	(5.222)	(0.377)
Loaned or swapped to other clubs before	0.343	0.296	0.047
	(0.475)	(0.457)	(0.035)
Average weekly wage before (£)	3,770.135	4,049.708	-279.573
	(3,384.244)	(3,631.801)	(293.912)
Maximum transfer fee before (£)	1,026,114.500	958,941.438	67,173.023
	(1,285,967.625)	(1,103,333.750)	(117,284.422)
# of players	364	372	

Notes: This table reports summary statistics for predetermined player characteristics for relegated and nonrelegated clubs. Columns (1) and (2) report the means at the relegation season, with standard deviations in parentheses. Column (3) reports the difference between relegated and nonrelegated players from a *t*-test, with standard errors in parentheses.

Table 2. Means of Appearance Rate, League Level and Wage by Relegation and Year

	Before relegation			First year after relegation			Five to seven years after relegation		
	Relegated	Non-relegated	Difference (1)-(2)	Relegated	Non-relegated	Difference (4)-(5)	Relegated	Non-relegated	Difference (7)-(8)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Appearance rate	0.509	0.494	0.015	0.543	0.473	0.070***	0.542	0.513	0.029*
	(0.345)	(0.349)	(0.014)	(0.313)	(0.307)	(0.026)	(0.316)	(0.322)	(0.017)
League Level	1.996	1.854	0.142	2.503	1.714	0.789***	2.951	3.495	-0.543***
	(2.989)	(2.051)	(0.109)	(2.802)	(1.232)	(0.181)	(2.716)	(4.385)	(0.202)
Wage (£)	4,350.786	4,377.327	-26.541	6,521.078	5,226.857	1,294.220	7,043.831	5,706.500	1,337.330**
	(4,578.530)	(5,249.323)	(309.086)	(10,903.070)	(6,156.492)	(960.341)	(11,938.255)	(7,864.206)	(634.099)

Notes: This table reports means for annual appearance rate, league level and wage for players in relegated and nonrelegated teams. Columns (1) and (2) report the means as of five years prior to the relegation season, Columns (4) and (5) report the means as of the first year after relegation, and Columns (7) and (8) report the means as of five to seven years after relegation. Standard deviations are in parentheses. Columns (3), (6) and (9) report the differences between relegated and nonrelegated players from a *t*-test, with standard error in parentheses. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table 3. Effect of Relegation on Annual Appearance Rate, League Level and Wage

Years since Relegation	**Dependent Variable**					
	Appearance Rate		League Level		Log (weekly wage)	
	(1)	(2)	(3)	(4)	(5)	(6)
Year 1	0.068**	0.059**	0.656***	0.732***	0.012	0.027
	(0.028)	(0.028)	(0.107)	(0.126)	(0.092)	(0.097)
Years 2-4	0.040	0.038	-0.247	-0.179	0.050	-0.003
	(0.025)	(0.026)	(0.185)	(0.210)	(0.103)	(0.104)
Years 5-7	0.033	0.031	-0.689**	-0.591**	0.232*	0.057
	(0.028)	(0.029)	(0.276)	(0.283)	(0.131)	(0.129)
Years 8-10	-0.016	-0.009	-0.599**	-0.486*	0.388**	0.237
	(0.033)	(0.034)	(0.289)	(0.287)	(0.187)	(0.159)
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
Season Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Player Fixed Effects		Yes		Yes		Yes
Observations	7,016	7,016	6,580	6,580	4,140	4,140
R-squared	0.108	0.371	0.080	0.363	0.113	0.670

Notes: Each row reports estimates for a dummy for being relegated interacted with relative-to-relegation year dummies. Values in parentheses are standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table 4. Heterogeneous Effect on League Level

Years since Relegation	Age		Prior Appearance	
	18 to 24	25 to 30	< Median	> Median
	(1)	(2)	(3)	(4)
Year 1	0.829***	0.648***	0.567**	0.793***
	(0.222)	(0.111)	(0.226)	(0.115)
Years 2-4	-0.021	-0.372	-0.315	-0.239
	(0.288)	(0.306)	(0.283)	(0.316)
Years 5-7	-0.862**	-0.303	-1.415***	-0.090
	(0.411)	(0.374)	(0.468)	(0.301)
Years 8-10	-0.699*	-0.128	-1.316***	-0.007
	(0.408)	(0.491)	(0.476)	(0.416)
Control Variables	Yes	Yes	Yes	Yes
Season Fixed Effects	Yes	Yes	Yes	Yes
Player Fixed Effects	Yes	Yes	Yes	Yes
Observations	3,656	2,924	3,387	3,193
R-squared	0.367	0.371	0.372	0.364

Notes: The dependent variable is a player's league level. Premier League is defined as level=1. Smaller values indicate better-ranked leagues. Each row reports estimates of a dummy for being relegated interacted with relative-to-relegation year dummies. Columns (1) and (2) use subsamples of players aged 18-24 and 25-30 at relegation season, respectively; Columns (3) and (4) use players appeared below and above the median at the relegation season, respectively. Values in parentheses are robust standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table 5. Heterogeneous Effect on ln (wage)

Years since Relegation	Age		Prior Appearance	
	18 to 24	25 to 30	< Median	> Median
	(1)	(2)	(3)	(4)
Year 1	0.009	0.012	0.068	0.066
	(0.151)	(0.129)	(0.131)	(0.137)
Years 2-4	0.020	0.047	0.134	0.026
	(0.157)	(0.140)	(0.149)	(0.128)
Years 5-7	0.308*	-0.137	0.406*	-0.065
	(0.186)	(0.182)	(0.207)	(0.169)
Years 8-10	0.525**	-0.065	0.693***	0.067
	(0.231)	(0.240)	(0.258)	(0.223)
Control Variables	Yes	Yes	Yes	Yes
Season Fixed Effects	Yes	Yes	Yes	Yes
Player Fixed Effects	Yes	Yes	Yes	Yes
Observations	2,313	1,827	2,054	2,086
R-squared	0.699	0.630	0.669	0.619

Notes: The dependent variable is log weekly wage (in British pounds). Each row reports estimates of a dummy for being relegated interacted with relative-to-relegation year dummies. Columns (1) and (2) use subsamples of players aged 18-24 and 25-30 at relegation season, respectively; Columns (3) and (4) use players appeared below and above the median at the relegation season, respectively. Values in parentheses are robust standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Appendix

Figure A1. Number of Teams Promoted, Retained, and Relegated by League Level

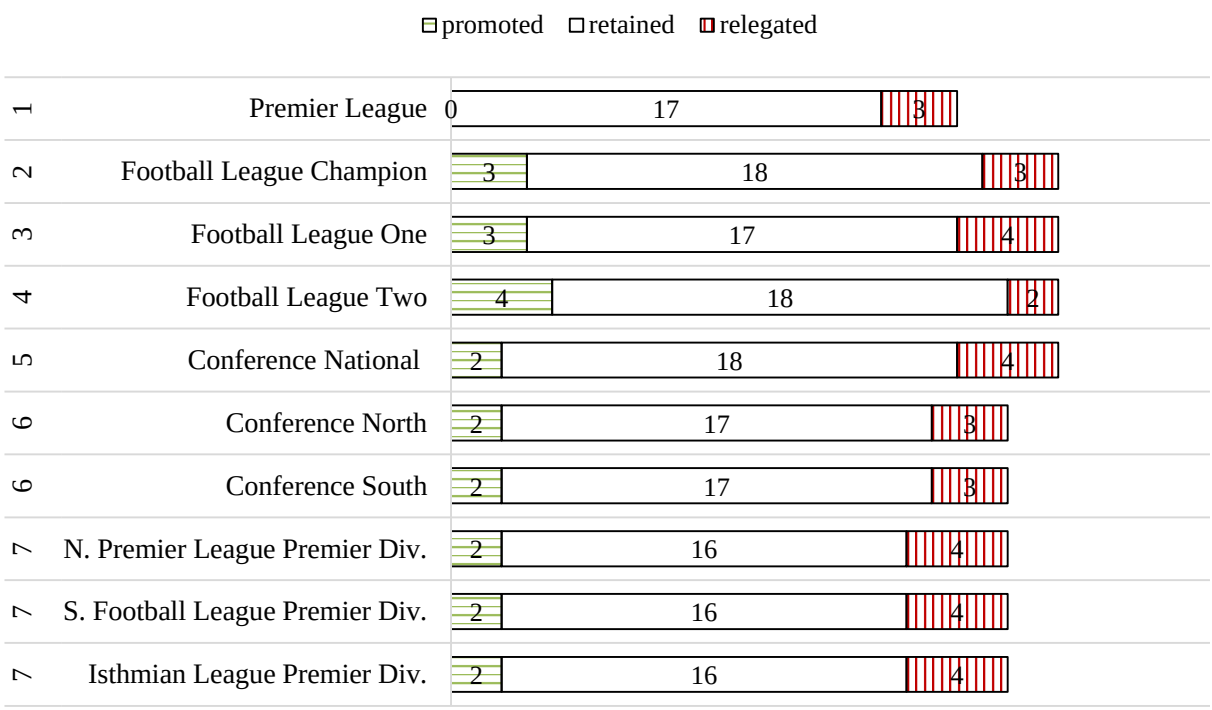

Source: English Football Association.

Notes: This figure shows the top 7 divisions of the English Soccer League and the number of teams promoted, retained, or relegated at each level.

Figure A2.1. Aggregate Revenue by League Level

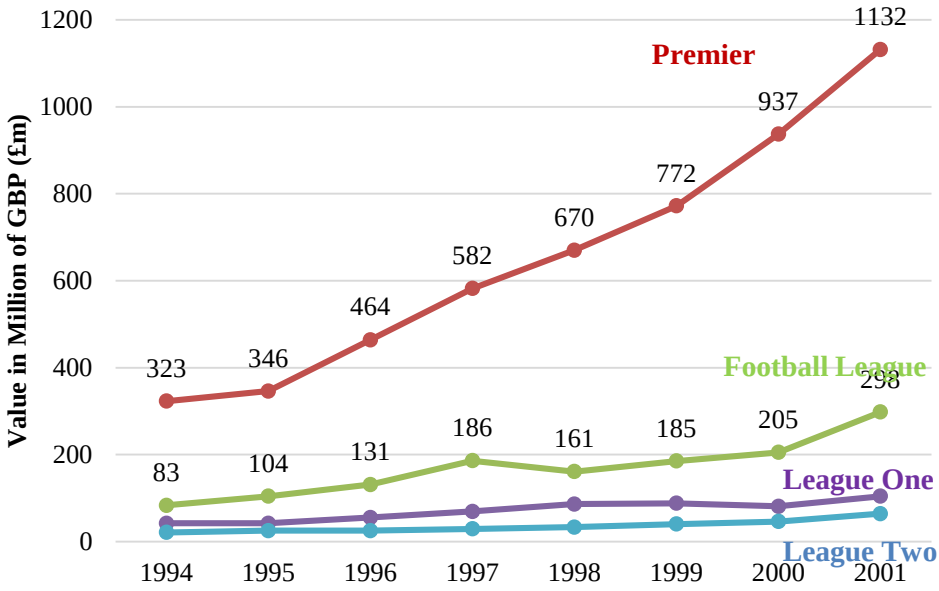

Figure A2.2. Average Annual Basic Pay of Players by League Level

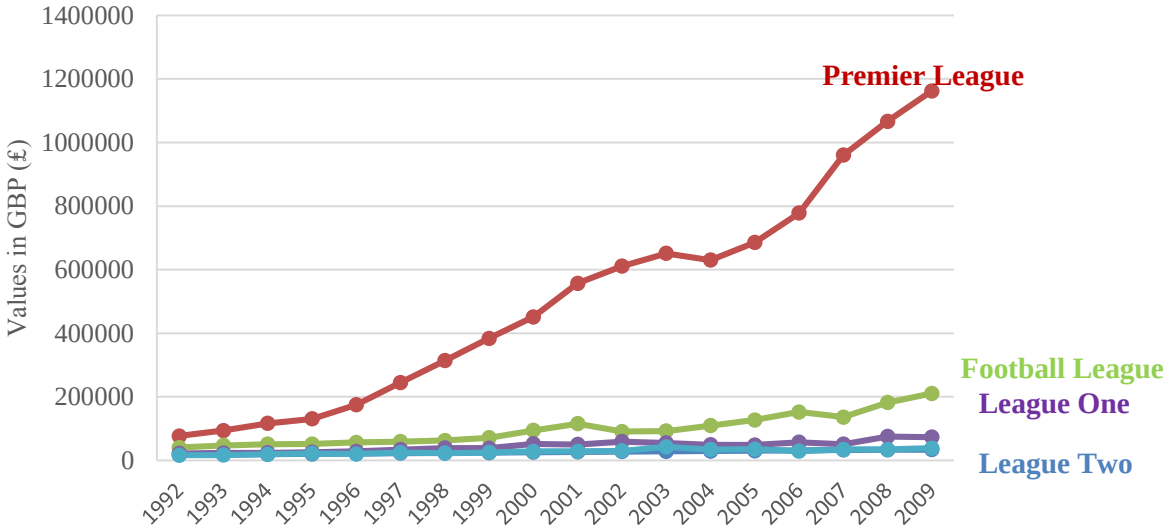

Notes: Aggregate revenue is the sum of revenues extracted from annual financial statements for each club at the same level in the league. Average basic pay does not include bonuses or commercial income.

Source: Figure A2.1: Annual Review of Football Finance, 2012, Deloitte. Figure A2.2: Professional Footballers' Association.

Figure A3. Percentage of Players Retained in Teams by Months after Relegation

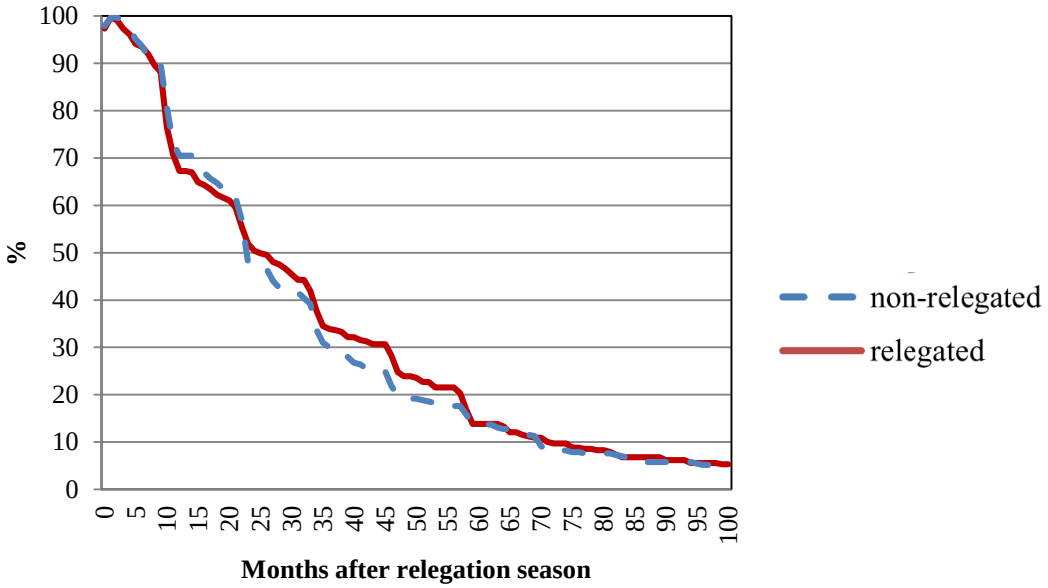

Table A1. Effect of Relegation on Player Dropout and Transfers

	Length of time from relegation to dropout	Length of time from relegation to transfer out
	(1)	(2)
Relegated	0.038	-0.053
	(0.033)	(0.331)
Player Characteristics	Yes	Yes
Relegation Season Fixed Effects	Yes	Yes
Observations	461	589
Pseudo R-squared	0.071	0.099

Notes: The table reports Poisson estimates on the length of time from relegation to dropping out of the sample (Column 1) and length of time from relegation to transferring out of the team (Column 2). Player characteristics include the age at relegation. Values in parentheses are robust standard errors. *, **, and *** represent statistical significance at a *p*-value of 0.10, 0.05 and 0.01, respectively.

Table A2. Characteristics of Dropouts and Transfers

Predetermined player characteristics	Relegated club players	Nonrelegated club players	Difference (1)-(2)
	(1)	(2)	(3)
Panel A. Players who dropped out within five years after relegation season			
Height (m)	1.803	1.820	-0.017
	(0.039)	(0.063)	(0.016)
Weight (KG)	75.541	78.721	-3.180
	(7.054)	(7.121)	(2.196)
Foreigner (Yes=1, No=0)	0.333	0.357	-0.024
	(0.482)	(0.488)	(0.135)
Age at relegation season	26.875	25.500	1.375
	(3.167)	(3.776)	(0.976)
Years of experience before relegation	8.875	7.321	1.554
	(3.405)	(4.355)	(1.098)
Loaned or swapped to other clubs before	0.478	0.261	0.217
	(0.511)	(0.449)	(0.142)
Average weekly wage before (£)	3,869.980	4,211.818	-341.838
	(3,406.422)	(3,082.075)	(989.784)
Maximum transfer fee before (£)	1,035,769.250	1,243,750.000	-207,980.766
	(971,236.000)	(1,112,436.125)	(416,796.938)
# of incidences of leaving	57	54	
Panel B. Players who transferred to other clubs within one year after relegation season			
Height (m)	1.809	1.805	0.004
	(0.071)	(0.069)	(0.019)
Weight (KG)	79.208	76.910	2.298
	(8.256)	(6.922)	(2.105)
Foreigner (Yes=1, No=0)	0.250	0.133	0.117
	(0.442)	(0.346)	(0.107)
Age at relegation season	25.042	24.200	0.842
	(3.196)	(3.478)	(0.919)
Years of experience before relegation	7.625	6.433	1.192
	(3.987)	(3.598)	(1.034)
Loaned or swapped to other clubs before	0.417	0.414	0.003
	(0.504)	(0.501)	(0.139)
Average weekly wage before (£)	4,484.777	3,791.994	692.783
	(4,182.799)	(2,297.018)	(1,011.216)
Maximum transfer fee before (£)	1,605,000.000	1,047,625.000	557,375.000
	(2,117,568.500)	(1,081,585.500)	(557,954.375)
# of incidences of leaving	36	36	

Notes: Panel A reports summary statistics for the characteristics of players who dropped out of professional leagues within five years after relegation. Panel B reports those of players who transferred out to other clubs within one year after relegation. Columns (1) and (2) report the means at the relegation season for relegated and nonrelegated players, respectively, with standard deviations in parentheses. Column (3) reports the difference between the two groups of players from a *t*-test, with standard errors in parentheses.

Table A3. Robustness Check: Exclude Transfers within One Year after Relegation Season

Years since Relegation	Dependent Variable					
	Appearance Rate		League Level		Log (weekly wage)	
	(1)	(2)	(3)	(4)	(5)	(6)
Year 1	0.059**	0.049*	0.681***	0.769***	0.000	0.019
	(0.029)	(0.029)	(0.115)	(0.137)	(0.095)	(0.106)
Years 2-4	0.024	0.020	-0.298	-0.225	0.037	0.010
	(0.025)	(0.026)	(0.203)	(0.231)	(0.106)	(0.106)
Years 5-7	0.041	0.038	-0.725**	-0.625**	0.232*	0.071
	(0.029)	(0.030)	(0.308)	(0.314)	(0.138)	(0.140)
Years 8-10	-0.024	-0.020	-0.500	-0.383	0.375*	0.225
	(0.033)	(0.035)	(0.319)	(0.321)	(0.198)	(0.170)
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
Season Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Player Fixed Effects		Yes		Yes		Yes
Observations	6,359	6,359	5,968	5,968	3,720	3,720
R-squared	0.118	0.385	0.075	0.362	0.121	0.675

Notes: The sample excludes the players who left the teams during the first season after relegation. Each row reports estimates for a dummy for being relegated interacted with relative-to-relegation year dummies. Values in parentheses are standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table A4. Robustness Check: Excluding Predictable Relegations

Years since Relegation	**Dependent Variable**					
	Appearance Rate		League Level		Log (weekly wage)	
	(1)	(2)	(3)	(4)	(5)	(6)
Year 1	0.072**	0.057*	0.696***	0.793***	-0.010	0.009
	(0.032)	(0.032)	(0.128)	(0.148)	(0.106)	(0.113)
Years 2-4	0.054*	0.044	-0.172	-0.113	0.046	-0.033
	(0.029)	(0.030)	(0.218)	(0.251)	(0.120)	(0.117)
Years 5-7	0.054*	0.050	-0.626**	-0.584*	0.142	0.008
	(0.031)	(0.033)	(0.295)	(0.308)	(0.150)	(0.147)
Years 8-10	0.005	0.011	-0.702**	-0.630*	0.298	0.180
	(0.038)	(0.040)	(0.356)	(0.349)	(0.220)	(0.183)
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
Season Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Player Fixed Effects		Yes		Yes		Yes
Observations	5,807	5,807	5,442	5,442	3,373	3,373
R-squared	0.113	0.380	0.077	0.366	0.120	0.691

Notes: The sample excludes the players in teams that were promoted to the Premier League, stayed for only one season, and were immediately demoted. Each row reports estimates for a dummy for being relegated interacted with relative-to-relegation year dummies. Values in parentheses are standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table A5. Effect of Relegation on Wage Attrition

Years since Relegation	Missing Wage Data	
	(1)	(2)
Year 1	-0.052	-0.054
	(0.034)	(0.035)
Years 2-4	0.014	0.011
	(0.029)	(0.031)
Years 5-7	-0.039	-0.047
	(0.033)	(0.034)
Years 8-10	-0.025	-0.044
	(0.041)	(0.042)
Control Variables	Yes	Yes
Season Fixed Effects	Yes	Yes
Player Fixed Effects		Yes
Observations	7,016	7,016
R-squared	0.364	0.511

Notes: The dependent variable is a dummy variable indicating missing wage data. Each row reports estimates for a dummy for being relegated interacted with relative-to-relegation year dummies. Values in parentheses are robust standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table A6. Effect of Relegation by Player Technique Index

Years since Relegation	League Level	
	< =Median	> Median
	(1)	(2)
Year 1	0.781***	0.548***
	(0.192)	(0.158)
Years 2-4	-0.062	-0.438
	(0.256)	(0.319)
Years 5-7	-1.001*	-0.598*
	(0.534)	(0.344)
Years 8-10	-0.101	-0.806**
	(0.569)	(0.337)
Control Variables	Yes	Yes
Season Fixed Effects	Yes	Yes
Player Fixed Effects	Yes	Yes
Observations	2,402	3,659
R-squared	0.383	0.390

Notes: The dependent variable is a player's league level. Each row reports estimates of a dummy for being relegated interacted with relative-to-relegation year dummies. Columns (1) and (2) use subsamples of players with a technique index below and above sample median at the relegation season, respectively. Values in parentheses are robust standard errors, clustered at player level. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.

Table A7. Heterogeneous Effects: Test Estimates between Subsamples

Years since Relegation	H_0: Equal effects for young and old		H_0: Equal effects by previous appearance	
	League Level (1)	Ln (wage) (2)	League Level (3)	Ln (wage) (4)
Year 1	0.181 (0.249)	-0.002 (0.201)	-0.226 (0.251)	0.003 (0.189)
Years 2-4	0.351 (0.401)	-0.027 (0.210)	-0.076 (0.420)	0.108 (0.183)
Years 5-7	-0.559 (0.539)	0.445* (0.259)	-1.326** (0.532)	0.471* (0.269)
Years 8-10	-0.571 (0.647)	0.590* (0.336)	-1.309** (0.645)	0.626* (0.348)

Notes: Each cell reports the difference in relegation's effects between two subgroups, with standard errors in the parenthesis. Columns (1) and (2) report the results of tests comparing younger (18-24 years) and older (25-30 years) players. Columns (3) and (4) report the results of tests comparing players appeared below and above the median at relegation season. *, **, and *** represent statistical significance at *p*-value of 0.10, 0.05 and 0.01, respectively.