

# The Age-Wage-Productivity Puzzle: Evidence from the Careers of Top Earners

by **Rachel Scarfe, Carl Singleton, Adesola Sunmoni  
and Paul Telemo**

[Discussion Paper](#) No. 2022-07

Department of Economics  
University of Reading  
Whiteknights  
Reading  
RG6 6EL  
United Kingdom

[www.reading.ac.uk](http://www.reading.ac.uk)

# The Age-Wage-Productivity Puzzle: Evidence from the Careers of Top Earners

Rachel Scarfe\*

University of Edinburgh

Carl Singleton<sup>†</sup>

University of Reading

Adesola Sunmoni<sup>‡</sup>

University of Reading

Paul Telemo<sup>§</sup>

University of Edinburgh

June 2023

[Link to latest version](#)

## Abstract

There is an inverted u-shaped relationship between age and wages in most labour markets and occupations, but the effects of age on productivity are often unclear. We use panel data on productivity and salaries in a market of high earners, professional footballers (soccer players) in North America, to estimate age-productivity and age-wage profiles. We find stark differences between these profiles; wages continue to increase for several years after productivity has peaked, before dropping sharply at the end of a career. This discrepancy poses the question: why are middle-aged workers seemingly overpaid relative to their contemporaneous productivity? The richness of our dataset allows us to investigate a range of possible mechanisms that could be responsible, including institutional factors, unobserved elements of productivity, and a talent discovery theory, by which teams pay younger players less because their productivity is more uncertain. We find some evidence that tentatively supports this last mechanism.

*Keywords:* Labour productivity; Wages; Ageing; Sports Labour Markets

*JEL codes:* J23, J24, J31, J41, Z22

---

\*School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK; [rscarfe@ed.ac.uk](mailto:rscarfe@ed.ac.uk) (corresponding author)

<sup>†</sup>Department of Economics, University of Reading, Edith Morley Building, Reading, RG6 6EL, UK; [c.a.singleton@reading.ac.uk](mailto:c.a.singleton@reading.ac.uk)

<sup>‡</sup>Department of Economics, University of Reading, Edith Morley Building, Reading, RG6 6EL, UK; [a.o.sunmoni@pgr.reading.ac.uk](mailto:a.o.sunmoni@pgr.reading.ac.uk)

<sup>§</sup>School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK; [paul.telemo@ed.ac.uk](mailto:paul.telemo@ed.ac.uk)

We are grateful for helpful comments and advice from participants in 2022 at the European Association of Labour Economists conference, the Scottish Economic Society Annual Conference (Glasgow), the XVI Gijón Conference on Sports Economics, and in 2023 at the 6th Erasmus Centre for Applied Sports Economics (ECASE) Workshop (Rotterdam).

# 1. Introduction

There is an inverted u-shaped relationship between age and wages in most labour markets and occupations: *ceteris paribus*, older workers earn more, until a peak is reached around age 50 (Mincer, 1958; Huggett et al., 2011; Rupert and Zanella, 2015). This age-wage profile is a key input in many theoretical labour market models, particularly those which feature worker choices and earnings over the life cycle. These models typically interpret the age-wage profile as an increase in productivity over time, either through investments in human capital or through returns to experience (e.g., Huggett et al. 2011). However, researchers can rarely observe a worker’s life-cycle or career-long productivity. They instead often rely on an underlying assumption, according to the human capital model of investment (e.g., Mincer, 1974; Brown, 1989) and proxying age for experience, that individual productivity is at least proportional to contemporaneous wages, implying that age-wage and age-productivity profiles should be similar. Our findings, from a labour market where both productivity and wages are directly observable, challenge that assumption.

It is now broadly accepted that there are many advantages to using data from sports labour markets to consider questions in labour economics (Kahn, 2000; Simmons, 2022). In our analysis, we use panel data for professional football (soccer) players in the United States and Canada. MLS is the premier association football league in North America, and our dataset covers approximately the universe of workers between 2008 and 2019. Media and fan interest in professional football ensures that detailed and accurate information on individual productivity over time is widely available. Another advantage of our setting is the League’s structure; MLS is a single-entity with a workers’ union, which negotiates salary regulation with MLS and publishes annual salary data for all the players (workers).<sup>1</sup> Within the rules set by MLS, which operates as a sort of cartel to compete with other professional sports in North America and other global football leagues, the teams or firms within the League are rivals in attracting and retaining players. The teams can therefore be thought of as price-takers for talent, even if the League’s actions may have some effects on that price.<sup>2</sup> Such effects may be limited, however, compared with the other professional sports cartels in North America, because of the fierce and relatively deregulated global competition for football talent (e.g., Frick, 2007a, 2009).

We find stark differences between the age-wage and age-productivity profiles in MLS; while productivity tends to peak at the age of 26 (approximately the middle of a footballer’s career), wages continue to increase into the early 30s (towards the end of their career). The richness of

---

<sup>1</sup>Due to this unique setting and data availability, MLS has been used previously to test theories of the labour market, for example: Coates et al. (2016) found a negative relationship between within-firm wage inequality and firm performance; Scarfe et al. (2020) showed that a team’s success on the pitch tends to correlate with consistently paying a relatively low price for talent; Scarfe et al. (2021) tested whether the productivity or popularity based theories of Rosen (1981) and Adler (1985), respectively, could best account for the superstar wages of top earners; and Butler and Coates (2022) showed that MLS goalkeepers generally receive a wage penalty, and outfield players who can perform multiple roles receive a premium.

<sup>2</sup>See Ferguson et al. (2000) for a theoretical and empirical analysis of the somewhat similar market for Major League Baseball talent.

our data allows us to estimate precisely an age-wage *premium* profile. We find that players at both the start and end of their careers are paid relatively less than their individual productivity would suggest compared with those in the middle of their careers. The magnitude of this age-wage premium is significant. Our estimates suggest that a 30-year-old footballer is paid roughly 40% more than a 20-year-old with the same level of performance and contribution to the League. We find some similar evidence in a major European football league, the German Bundesliga, though using less reliable salary data.

We investigate three potential reasons for the age-wage *premium* profile among professional footballers: institutional factors specific to this labour market; unobserved productivity that varies with age but is not captured in our data; and market imperfections that mean a worker is not paid their marginal productivity. Our data allow us to test these explanations, and in Section 5, we discuss them in more detail. First, we consider the role of salary regulations or collective bargaining, as previous research suggests that wages increase more with seniority in more unionised industries (De Hek and van Vuuren, 2011; Williams, 2009; Zangelidis, 2008). The collective bargaining agreements (CBAs) between the employer, Major League Soccer (MLS), and the union, the MLS Players Association, are publicly available. Although these agreements specify minimum salaries and annual increases, they affected very few players. We find that the age-wage profile estimates are approximately unchanged by the CBAs despite them being renegotiated twice in our study period. This suggests that the estimated profile of the age-wage premium is not sensitive to the institutional factors unique to MLS.

We also test whether there are elements of productivity that correlate with age but are not easily observed. Some talent may become increasingly popular as it ages, attracting more fans and revenue to stadiums and the League as a whole. There is evidence that such a mechanism plays a role in wage setting among high earners in a range of industries (Carrieri et al., 2018; Filimon et al., 2011; Hoffman and Opitz, 2019). A minority of workers in MLS are ‘designated players’. These tend to be more the popular and higher ability players (Kuethe and Motamed, 2010), specifically recruited with the intention of increasing the popularity of the League, and are subject to different salary regulation (Coates et al., 2016). However, we find that dropping these players from our analysis does not change the shape of the estimated age-wage premium profile. Further, we use page views of player Wikipedia profiles as a proxy for popularity, finding that estimated *age-popularity* profiles correspond closer to our age-productivity profiles than to age-wage profiles. This provides further evidence that superstar effects cannot explain the career wage-productivity discrepancies that we observe. We also ask whether mid-career workers may have accumulated human capital that is not captured by ‘direct’ productivity measures, such as leadership skills, by comparing the outcomes of teams that chose different age distributions of employees in their workforce. We use the variation in the age distribution of players across teams, to test whether teams who pick more players with high age-wage premiums win more football games. We find no evidence that this is the case. After controlling for the average wage rate of a team, picking players with high age-wage premiums (i.e., around age 30) is associated with winning fewer football games in MLS.

Finally, we consider the role of market imperfections in generating the age-wage premium. We begin by testing a talent discovery mechanism, by which firms pay younger workers less because their productivity is less well-known. In our setting, if a young player has a particularly good season, teams are unable to distinguish whether this is evidence of permanently high productivity or due somewhat to luck or other temporary factors. This should be less of an issue for older workers, for whom more past performance data is available. This mechanism is similar to that described by [Terviö \(2009\)](#), whereby workers and firms are not able to commit to long-term wage contracts and productivity is only revealed through actual on-the-job performance. In such labour markets, [Terviö](#) showed that firms will excessively bid for known talent rather than trying out new talent. To test this, we regress a player's latest productivity on that observed during previous seasons of the League. We find some, limited, evidence that more observations of past productivity helps to predict future productivity for relatively young players. This suggests that teams could favour players with more experience in the League, explaining some of the observed age-wage premium, although further investigation is needed. But in general, none of these explanations we have considered appears to account substantively or conclusively for the seemingly ungenerous wages that the youngest and oldest talent tends to earn in MLS. This leaves us with an unsolved puzzle: why are some workers paid so handsomely as they move toward the mid-point and later stages of their careers despite their dwindling talent?

There are a number of previous studies estimating the age-productivity profiles of workers. In general, these studies find that productivity increases initially during a career, before declining or remaining flat. The combination of individual productivity and salary data covering all workers in a labour market is rare. Previous studies have looked at settings where individual productivity is measurable, including sports, to estimate age-productivity profiles without comparing these to wage profiles (e.g., [Fair and Kaplan, 2018](#); [Fair, 2008](#); [Oster and Hamermesh, 1998](#); [Börsch-Supan et al., 2021](#)), or used firm-level rather than individual worker-level productivity data (e.g., [Cardoso et al., 2011](#); [Hellerstein et al., 1999](#); [Van Ours and Stoeldraijer, 2011](#)). Our data allows us to estimate profiles that are robust to selection in and out of a labour market, and to both unobserved differences in workers' average productivity and intangible worker-firm-specific factors that are valued by employers. We show that failing to account for these effects leads to substantially biased estimates of the age-productivity and age-wage profiles. This echoes the findings with respect to productivity of [Bertoni et al. \(2015\)](#), [Castellucci et al. \(2011\)](#) and [Hakes and Turner \(2011\)](#), for chess players, Formula One racing drivers and Major League Baseball players, respectively. Our setting has a further advantage: football players generally perform the same tasks (training, playing matches, making promotional appearances, etc.), albeit with variation across playing positions, and are all on full-time contracts. This is not necessarily the case in other settings, where managers may assign different tasks to older or younger workers, even with the same job title, and hours of work can be poorly measured or subjectively reported by employees, which normally make estimation of the effects of age on productivity challenging.

Although we study a particular labour market, there is some, currently limited, evidence that an age-wage premium exists more widely. [Cardoso et al. \(2011\)](#) analysed administrative data

covering all private manufacturing and services firms in Portugal, finding that workers close to the ends of their careers are relatively underpaid compared with their contributions to firm-level productivity. In contrast, [Ilmakunnas and Maliranta \(2005\)](#) and [Dostie \(2011\)](#) found wage premiums for older workers using data covering all Finnish manufacturing firms and all Canadian workplaces, respectively. [Hellerstein et al. \(1999\)](#) and [Van Ours and Stoeldraijer \(2011\)](#) did not find evidence that older workers are overpaid, but did not consider whether younger workers are underpaid. These studies all used firm-level rather than individual worker-level productivity data. Our findings add to the evidence in support of a wage premium for mid-career workers, by estimating an individual's age-wage premium profile. Importantly, we provide evidence for a premium in a market in which careers are short and without long-term contracting, and in which effort and performance are easy to monitor. Thus, the usual theories of employers using delayed compensation to incentivise effort and loyalty ([Lazear, 1979, 1981](#)) are unlikely to apply in this setting, unless both the transaction costs of renegotiation or piecemeal contracting are very high and the workers are sufficiently risk averse. However, this still appears to be contradicted by the relatively short tenures and high mobility seen in the football talent market. Studying a small sample of Major League Baseball Players and perhaps most closely related to our own analysis, [Blass \(1992\)](#) concluded that wages tended to increase with experience irrespective of productivity gains. [Blass](#) also agreed with ourselves that it was not obvious why implicit contracts might be needed in such a market, given that productivity is so clearly observed.

Using our individual productivity data and exploiting variation in the distribution of ages across different teams, we are able to test a number of other potential explanations for the age-wage premium, as discussed above, and conclude that it is unlikely to be explained by institutional factors, or by unobserved elements of productivity. We find some evidence aligning with a theory that lack of information on the productivity of younger workers can contribute to a wage premium that rises with age. In most settings it is impossible to observe the relevant information that managers have when hiring a worker, but in sports labour markets we can observe past performance. There are some previous research findings that firms prefer to hire experienced workers, even if they are less productive in expectation ([Peeters et al., 2022](#) on football managers; [Pallais, 2014](#) in a field experiment; and [Stanton and Thomas, 2016](#) on the value of intermediaries in online hiring for young workers). Our findings add to this body of evidence, which suggests that firms in some labour markets may be inefficiently and excessively sampling from the pool of known or experienced talent, at the expense of discovering and developing new talent.

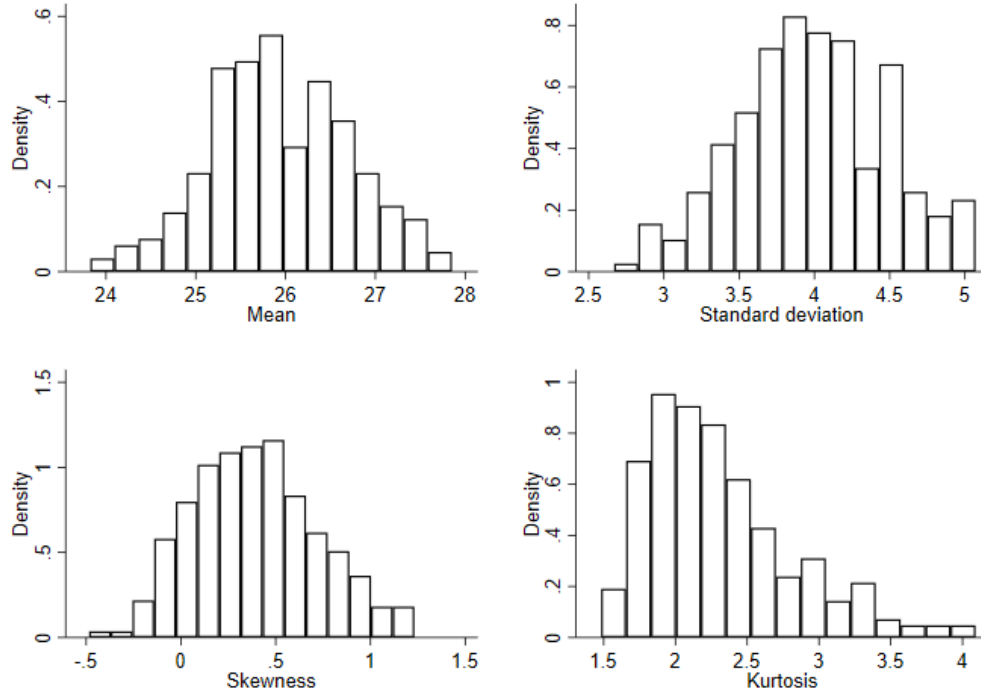
## 2. Setting and data

### 2.1 Worker ages

We obtained data on worker ages from the [MLS official website](#). Figure 1 shows the distributions over every firm (MLS team) and year combination (245 in total) in our dataset for the mean, standard deviation, skewness and kurtosis of ages on rosters. Figure 2 shows the the age profile

of workers over time, at three-yearly intervals. This has been broadly stable, although there were more older workers later in the period that we study. In general, as Figures 1 and 2 show, the distribution of age within a firm tends to be positively skewed, with a large number of players in their early 20s and fewer in their late 20s and early 30s. The kurtosis is also generally small, with more workers in the centre of the age distribution and fewer very old or young workers.

FIGURE 1: Distributions over firm-years for statistics of worker ages on rosters, 2007-2019




## 2.2 Wages

The wage data published by the MLS Players Association (MLSPA) capture the mid-season of the MLS in August, after the secondary transfer window when players can be signed from other leagues, and cover the 2007–2019 seasons. The measure of wages that we use is the guaranteed annualised compensation, henceforth referred to just as wages. This is a worker’s total base annual salary over the years covered by their contract, plus payments for signing with a team or related to marketing, divided by the number of years covered by the contract. It does not include performance related payments. However, MLS’s salary regulations require that any “readily achievable” individual bonuses are reflected in the guaranteed annualised compensation published by the MLSPA. Our measure of wages is therefore little affected if a player has a particularly good or bad year and thus receives performance bonuses that are much higher or lower than expected.

To account for general increases in wages over time, both due to inflation and to changes in the MLS CBAs, we detrend by season and normalise. We first subtract from each wage observation the mean wage in the sample for each respective season. We then take these residual or detrended wages and add the average wage in 2019. The wages we analyse can, therefore, be thought



FIGURE 2: Age distribution of workers, selected years in 2007-2019



Notes.- Kernel density estimated using a Gaussian kernel and automatic plugin bandwidth selection.

of as being in “MLS 2019 prices”, accounting for general League trends in average pay. This de-trending approach also simplifies our age profile estimation and interpretation.<sup>3</sup>

Figure 3 displays the distribution of year-to-year changes in the natural logarithm of our wage variable. This variation is particularly relevant as we later use worker-level fixed effects to account for the selection of workers in and out of MLS. The majority of annual wage changes are small. However, there are also fairly frequent cases of large raises or cuts in annual salaries.


## 2.3 Individual productivity and popularity

Our key productivity measures are average season-level player ratings, which we obtain from [WhoScored.com](https://www.whoscored.com), available from 2013, and minutes played in a regular season. [WhoScored.com](https://www.whoscored.com) constructs player ratings for matches in top football competitions around the world using data from [Opta](https://www.optasports.com), a market-leading British sports analytics company that provides raw data for thirty different sports in seventy countries and is the official supplier of statistics to MLS and media organisations, including the British Broadcasting Corporation (BBC) and Sky Sports in the UK. To generate ratings on a scale of 1–10 for every player live during a football match and over its duration, [WhoScored.com](https://www.whoscored.com) uses a unique and comprehensive statistical algorithm. Thus the

<sup>3</sup>In a linear regression model of wages with worker fixed effects, as well as dummies or linear terms for age and time, the marginal effects of an individual ageing are not constant but instead dependent on the time period - both age and time increase by one unit for a player between season. In other words, the age and year effects cannot both be identified. This is the classic problem of distinguishing age, year and birth cohort effects in linear regression analysis, where the latter can be loaded into worker fixed effects. Residualising and normalising wages (and productivity) over time, using a first-stage regression model or just de-meaning the data, allows for a flexible and robust estimation of age profiles (see Section 3B of [Card et al. \(2018\)](#) for a recent discussion on this topic, or [Hanoch and Honig \(1985\)](#) for an early application of this approach to solving the problem).



FIGURE 3: Distribution of year-to-year changes in wages, 2007-2019



Notes.- this figure excludes the top and bottom 1% of observations for clarity, although they are included in the main estimation samples

. Wages are always expressed and analysed in “MLS 2019 prices” (see text). Sample size of annual wage changes over this period is 3,195.

ratings are likely to be freer of bias than more subjective ratings, such as those produced by journalists ([Principe and van Ours, 2022](#)). Over 200 raw statistics are included in the computation of a player’s rating, weighted by their influence within the game. All events of importance are taken into consideration, with a positive or negative effect on ratings weighted in relation to the area on the football pitch and the outcome. For example, an attempted dribble in the opposing team’s final third that is successful will have a positive effect on a player’s rating. According to [WhoScored.com](#), ratings less than 5.9 are “poor”, ratings of 6.0 – 6.9 are “average”, 7.0 – 7.9 are “good”, 8.0 - 8.9 are “very good” and 9.0 - 10 are “excellent”.<sup>4</sup> [Holmes and McHale \(2023\)](#) have shown that these ratings can be used to effectively forecast football match results within a dynamic player-based model, performing well against bookmaker odds-implied forecasts.

We use the [WhoScored.com](#) ratings data from the 2013-2019 seasons of MLS, which we merge with the other data described above, as well as for the top division of German professional football (Bundesliga) over the same period for a later robustness check. These player ratings provide an alternative measure of productivity, capturing a player’s overall performance in a way that is relevant and comparable for all positions (forwards, midfielders, defenders and goalkeepers). They are widely used by the media and bookmakers.<sup>5</sup> As we are interested in the relative ratings between players, we detrend or residualise them by season for each position, to account for any

<sup>4</sup>See [whoscored.com/Explanations](#) for a more detailed description of how these ratings are calculated.

<sup>5</sup>For example, [WhoScored.com](#) provides the ratings for the “Man of the Match” betting markets for the major bookmakers Betfair and Betvictor and for the Sun newspaper’s fantasy football website.

changes over time in the algorithm or information set used by [WhoScored.com](#), using a similar method to the one described above for wages.

Minutes played is another suitable proxy for individual productivity, as better players will tend to be on the pitch for more minutes over a season, assuming that football managers are aiming to win matches.<sup>6</sup> Success on the pitch is also positively associated with financial returns in most major North American sports leagues, including MLS ([Bradbury, 2019, 2021](#)). Further, team owners may want players who are more popular with fans to play more. If team owners and the managers they choose are revenue maximisers, and play more popular players ahead of those with higher on-pitch performance, then this will be reflected in minutes played. Another advantage of this productivity measure is that it takes account of injuries or suspensions for foul play. These are not necessarily reflected in the [WhoScored.com](#) ratings, which only consider the time a player is on a pitch. A player that is injured is extremely unproductive, in the same way that workers in other labour markets who suffer health issues are unproductive. Thus the two measures of productivity that we use complement one another. Finally, we use minutes played as other readily available individual productivity measures, which we collected from the [MLS official website](#), such as goals, shots or assists per game, are more applicable to forwards and mid-fielders than defenders and goalkeepers.

We also use data on players' profile page views from [Wikipedia](#), which we use to proxy for their popularity, another potential element of individual worker productivity in MLS. As in [Scarfe et al. \(2021\)](#), we merge our sources of data using worker names and years, creating a dataset with 6,135 worker-year observations over 1,885 individual workers from 2007 to 2019. We drop a tiny number of observations due to missing age or productivity indicators, or because they could not be matched across our other data sources. Further, we exclude observations younger than age 19 or older than 35, since there are not enough observations to robustly estimate age profiles of wages or productivity at these extremes of careers in MLS.

## 2.4 Descriptive statistics

Table 1 presents some descriptive statistics for MLS players in 2007-2019. The average worker age is 25.9, with a standard deviation of 4.0 years. Minutes played per game within a season range from zero (for the 11% of players who acted as reserves for the whole season or who were injured) and 90 (for the 1% who played every minute of every game), with an average of 34.5. Since we later use the panel structure of the dataset, we also show descriptive statistics in Table 1 for workers who spend at least two years in MLS, which covers 58% of our sample.

---

<sup>6</sup>[Késenne \(2006\)](#) provides a detailed discussion of the incentives behind decisions that employers make in sports talent markets. Also, see [Burguet and Sákovics \(2019\)](#) for a theoretical discussion of how to accommodate simultaneously both the win and profit maximising objectives of sports team owners.

TABLE 1: Summary statistics over worker-year observations, 2007-2019

Variable	Mean	St. dev.	Min.	Median	Max.	N
<i>All Players</i>						
Log guaranteed salary (MLS 2019 prices)	12.27	0.934	10.7	12.2	17.1	6,135
Mins played per game	34.49	28.19	0.0	31.0	90.0	5,516
WhoScored ratings (1-10)	6.753	0.346	4.7	6.8	8.9	2,928
Age (years)	25.93	4.004	19.0	25.0	35.0	6,135
Tenure in MLS (years, continuous spell)	2.887	2.255	1.0	2.0	13.0	6,135
<i>Players with at least two seasons in MLS</i>						
Log guaranteed salary (MLS 2019 prices)	11.80	0.946	9.5	11.7	16.0	3,925
Mins played per game	39.19	28.30	0.0	39.3	90.0	3,513
WhoScored ratings	6.763	0.341	4.7	6.8	8.9	2,130
Age	26.54	3.957	19.0	26.0	35.0	3,925
Tenure	3.950	2.194	2.0	3.0	13.0	3,925

Notes.- Data on WhoScored ratings are from [WhoScored.com](https://www.whoscored.com) from 2013 onwards, data on guaranteed salaries are from the MLS Players Association and all other data are from the [the official MLS website](https://www.mlssoccer.com). Minutes played per game data is truncated at 90 minutes (the length of a full game). We use minutes per game from 2007 to 2018.

## 2.5 Institutional setting


Unlike most other top football leagues in Europe and around the world, MLS is closed and does not feature promotion or relegation. Changes in the composition of the League occur through franchise (team) expansion or dissolution. Teams compete in two parallel leagues: the Eastern and Western Conferences. Each team plays every other team in their conference twice each season (calendar year) and several games against teams in the other conference, so that each team plays thirty-four games in a regular season. The top six teams in each conference advance to the MLS playoffs, which is a knockout series to determine the championship winner, known as the MLS Cup. Up to five teams (four from the US and 1 from Canada) qualify each season to participate in an international competition, the CONCACAF Champions Cup, against other North American teams.

MLS is a single corporate entity and owns a stake in all the franchise teams, which receive some revenues directly, such as ticket sales, all stadium revenue, and local broadcast rights (Peeters, 2015). In addition, the teams receive a portion of the overall League's profits, including national and international broadcast rights, as well as sponsorship money (Scarfe et al., 2021). A consequence of this structure is a single union, the [MLS Players Association \(MLSPA\)](https://www.mlspa.com), which negotiates salary regulation with the League, resulting in a new Collective Bargaining Agreement (CBA) between the MLSPA and MLS every five years. The regulation governs all players in MLS and salaries are published each season by the MLSPA.

## 2.6 Contracts in MLS

The majority of players negotiate and sign a contract known as a Standard Player Agreement (SPA) with the League. SPAs are generally short, with an initial ‘guaranteed’ period of one or two years, plus up to three option years. After the initial term of an SPA has finished, teams can exercise their option to retain a player. Tenure in MLS is generally short (2.8 years on average), as shown in Figure 4. A professional footballer’s career is generally shorter than in other labour markets, with most entering in their late teens and retiring in their early 30s (Barth et al., 2021). Therefore, average tenure in MLS (2.8 years) is approximately a fifth of the average total career duration of a footballer. Job tenure in other labour markets varies across countries and sectors, ranging from estimates for the US, where workers hold an average of 12 jobs over their career (Bureau of Labor Statistics, 2021), to Germany, where workers hold an average of four jobs in the first 20 years of their careers (Dustmann and Pereira, 2008).

FIGURE 4: Distribution of worker tenure, 2007-2019



Notes.- Shows the lengths of 6,135 players’ longest continuous spells (seasons/years) in MLS.

Teams put together rosters of up to thirty players (as of the 2022 season) who can play for them in a season. Players can enter MLS through three processes: a ‘draft’ of junior players; a ‘discovery’ process, whereby teams can scout for players from other leagues; and an ‘allocation’ list maintained by the League of players who previously played in MLS before moving to another league, probably overseas. Teams can acquire players for their roster via these methods; through trades with other teams; and also through a ‘re-entry’ draft of players whose options have not been exercised by their current MLS team (Major League Soccer, 2021). There is also a limited form of ‘free-agency’, introduced in 2015, whereby older players with a relatively long tenure in MLS can negotiate directly with any team when their contracts expire. There are further regulations regarding which players a team can sign. These include limits on the number of international

players and a quota for the number of young players (aged under 24) that a team must include on its roster.<sup>7</sup> Consequently, in our data we observe players at all points in their careers.

MLS argues that this regulation ensures the league remains financially viable and improves competitive balance. This appears to have been successful, compared with major European football leagues, as the MLS Cup was won by twelve different teams in the thirteen seasons between 2007 and 2018 that we study.<sup>8</sup> This is an advantage from our perspective. Team managers can theoretically improve their team's performance in the League dramatically and in relatively quick time through their roster choices; if winning is an objective of the decision makers in a team, then they should be aiming to secure the most productive players at the lowest possible wages.

As we explain above, players negotiate their wages and sign contracts with the League, which acts as single-entity, rather than with individual teams. This raises concerns that our conclusions may lack external validity, even within the context of professional football. If MLS is indeed a near-perfect monopsony labour market, then this will affect salaries and mobility between teams negatively, and perhaps differently over the duration of a career, as discussed by [Ferguson et al. \(2000\)](#), [Kessenne \(2015\)](#) and [Scully \(1974\)](#) with regards the exploitation of talent within North American sports. However, although it operates as a single entity cartel, and in its early years kept salary costs low relative to revenue ([Twomey and Monks, 2011](#)), MLS must compete with other football leagues to attract and retain talent. It was for this reason that the designated player rule was introduced, to allow teams to spend above their salary cap on elite talent, typically recruited from outside the US ([Coates et al., 2016](#)). Further, a substantial proportion of players in MLS were born outside North America (see Online Appendix Figure [A10](#)), and many well-known American players have been successful in European Leagues. This suggests that MLS operates in monopsonistic competition with teams in other football leagues abroad, some of which have much greater market power, rather than being a near-perfect monopsonistic employer of talent, as is the case with other American sports leagues such as the National Football League.<sup>9</sup> As [Bradbury \(2021\)](#) notes, teams in MLS retain revenue from ticket sales and local broadcasts, and from the sale of players to teams in other football leagues, and there is a positive correlation between the player quality in MLS teams and their revenue. Despite any effects that MLS may have on the price for talent within the League, the teams regardless compete with each other to recruit and retain players. In this respect, we would argue that MLS is similar to other labour markets, in football at least, whether viewed either as a single employer, in competition with the rest of the world, or as multiple teams competing within MLS and the rest of the world subject to the particular constraints and regulation imposed by the League.

---

<sup>7</sup>See, for example, [Major League Soccer \(2020\)](#), and further discussion in Section [5](#) for more detail.

<sup>8</sup>For example, the Italian Serie A and German Bundesliga were both won by only three different teams over the same time period. Although this comparison is imperfect, since the European leagues do not have a playoff structure with finals and semi-finals to determine the eventual champion, it illustrates a difference in competitive balance which applies across a number of measures ([Prockl and Semmelroth, 2018](#)).

<sup>9</sup>The Bosman ruling of 1995 created something akin to a single talent market in European football, with spillover effects on other leagues worldwide. See [Ericson \(2000\)](#); [Fort \(2000\)](#); [Frick \(2007b, 2009\)](#); [Goddard et al. \(2012\)](#) for more discussion on the economics of the Bosman ruling and especially its effects on the international mobility of talent.

### 3. Age-productivity profiles

#### 3.1 Empirical strategy

We use least squares to estimate age-productivity profiles for a range of model specifications. As dependent variables, we use two measures of individual productivity for a worker over a season: average minutes played per game and their average [WhoScored.com](#) rating. We estimate the age-productivity profiles using both non-parametric models, with dummy variables for each year of age, and parametric models, with polynomials in age. In the first instance, we address the selection of workers in and out of this labour market over time using worker fixed effects. We also investigate how allowing for firm and worker-firm fixed effects in the models affect our results:

$$y_{i,t} = \alpha + \beta \mathbf{Age}_{i,t} + \varepsilon_{i,t} \quad (\text{OLS}) \quad (1)$$

$$y_{i,t} = \alpha + \beta \mathbf{Age}_{i,t} + \lambda_i + \varepsilon_{i,t} \quad (\text{Worker FE}) \quad (2)$$

$$y_{i,t} = \alpha + \beta \mathbf{Age}_{i,t} + \lambda_i + \phi_{J(i,t)} + \varepsilon_{i,t} \quad (\text{Worker and firm FE}) \quad (3)$$

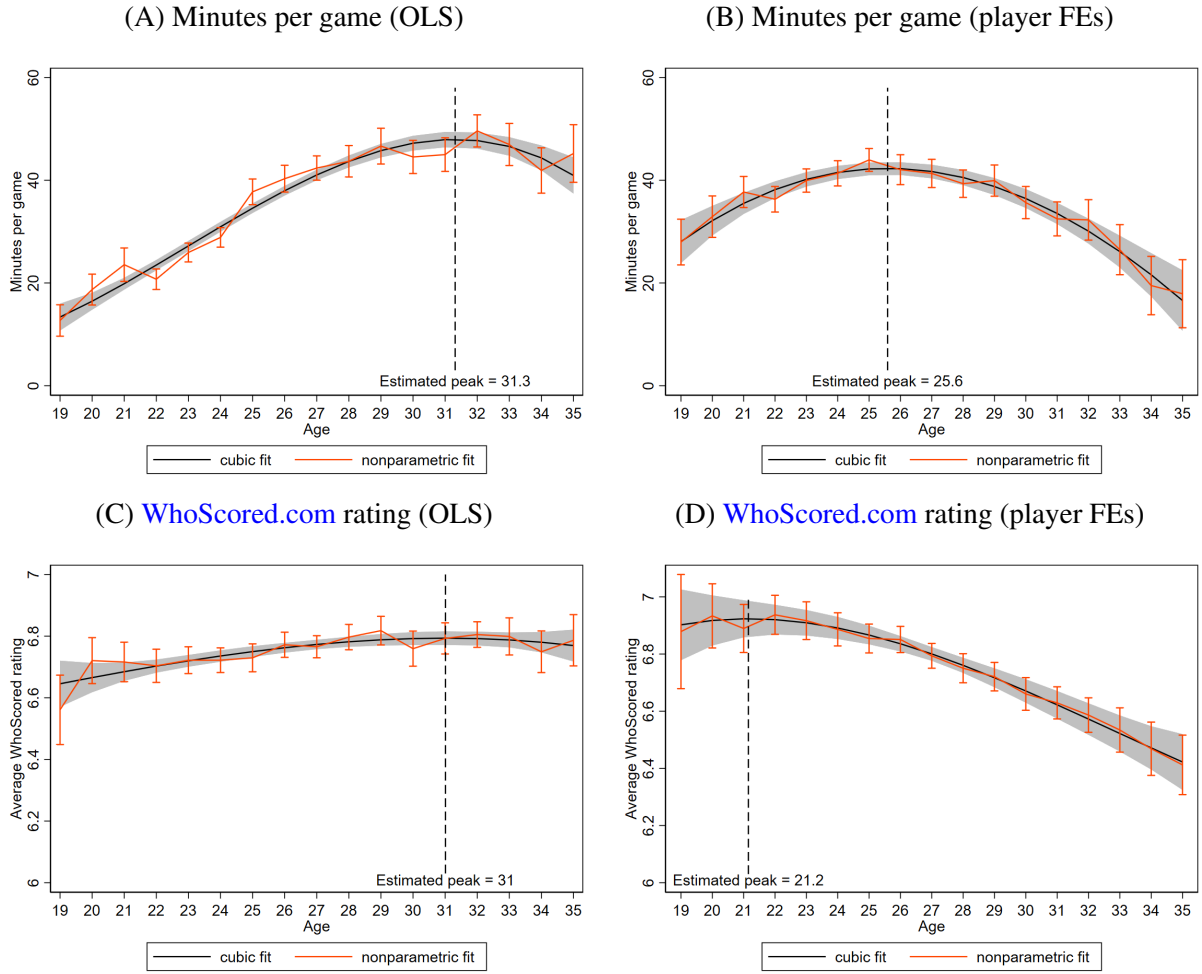
$$y_{i,t} = \alpha + \beta \mathbf{Age}_{i,t} + \gamma_{i,J(i,t)} + \varepsilon_{i,t} \quad (\text{Worker-firm FE}) \quad (4)$$

Subscript  $i$  denotes an individual worker,  $t \in \{2007, 2009, \dots, 2019\}$  denotes the year, and  $j = J(i, t)$  is the firm of individual  $i$  in year  $t$ . In the parametric models,  $\beta$  is the vector  $[\beta_1 \ \beta_2 \ \beta_3]$  and  $\mathbf{Age}_{i,t}$  denotes the vector  $[Age_{i,t} \ Age_{i,t}^2 \ Age_{i,t}^3]'$ , whereas for the non-parametric models  $\beta = [\beta_{20} \ \beta_{21} \ \dots \ \beta_{35}]$  and  $\mathbf{Age}_{i,t} = [\mathbb{1}\{Age_{i,t} = 20\} \ \mathbb{1}\{Age_{i,t} = 21\} \ \dots \ \mathbb{1}\{Age_{i,t} = 35\}]'$ , where age 19 is the excluded category. As explained in Section 2, we account for the possibility that the [WhoScored.com](#) system has changed over the years by normalising the ratings by year and position-specific means. We also report similar estimates using the raw ratings in Online Appendix Figure A4.

#### 3.2 Results

Figures 5A and 5C show the ‘naive’ OLS model estimates of footballer productivity peaking at age 31, both in terms of minutes played and [WhoScored.com](#) ratings. On average, the minutes per game ([WhoScored.com](#) ratings) increase from around 13 (6.6) at age 19 to 48 (6.8) by the peak at age 31. Accounting for worker-level fixed effects lowers these peak age estimates to 26 and 21 for minutes played and [WhoScored.com](#) ratings, respectively (Figures 5B and 5D). This suggests that weaker workers tend to drop out of this labour market as they age. Another possible explanation for these results is MLS’s history of purchasing older ‘superstar’ players from Europe, who, unless accounted for in a fixed effects model, will bias the age-productivity profile estimates upward for older players. Focusing on minutes played per game for a given worker, the peak estimate at age 26 is 43 minutes, compared with 28 minutes at age 19, and with a substantial decline towards the twilight of a career to 18 minutes by age 35.

FIGURE 5: Estimated *age-productivity* profiles using OLS and worker-level fixed effects for Major League Soccer



Notes.- Estimates of Equations (1) and (2): minutes per game concern 2007-18; WhoScored.com ratings concern 2013-2019. The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

We also consider whether firms of different quality pursue different recruitment strategies with respect to the ages of their workers. The worker and firm fixed effects, Equation (3), and worker-firm match fixed effects, Equation (4), model estimates do not differ in any meaningful way from those only addressing worker-level fixed effects, with peak ages remaining within one year of those from the latter more parsimonious specifications (Online Appendix Figure A1). This may be surprising at first glance. However, average minutes per game is bounded above by 90 and below by zero, and so is unlikely to differ greatly between firms. As discussed in Section 2, MLS also has stringent salary regulation, including a cap, which is aimed at maintaining both a competitive league and the financial stability of the franchise teams. As a result, the distribution of talent across teams is likely to be more even than in other football leagues, such that the team a player joins is less likely to affect their productivity. For this reason, we consider the more parsimonious worker fixed effects model as our preferred specification.

We report estimation results for the parametric specification of Equation 2 in Table 2. We solve for the estimated peak ages of productivity analytically and generate their 95% confidence



intervals using bootstrapping, which are 24.8-26.4 years for minutes per game and 17.2-25.1 for [WhoScored.com](#) ratings. In Online Appendix Figures A2 and A3, we also investigate whether the age-productivity profiles differ depending on playing positions, by estimating the worker-level fixed effects models separately depending on a player's main position (goalkeeper, defender, midfielder or forward). For minutes per game, we estimate that goalkeepers peak at age 31, whereas all other positions have estimated peaks at age 25. For the [WhoScored.com](#) ratings, we cannot estimate any significant age-productivity curve for goalkeepers. For defenders, we estimate a peak rating at age 20, for midfielders at 22, and for forwards at age 23.

TABLE 2: Estimated *age-productivity* profiles using OLS and worker-level fixed effects for MLS players, 2013-2019

	Minutes per game (I)	<a href="#">WhoScored.com</a> rating (II)
Age ( $\beta_1$ )	20.59 (12.96)	0.347 (0.247)
Age <sup>2</sup> ( $\beta_2 \times 10$ )	-4.919 (4.876)	-0.121 (0.091)
Age <sup>3</sup> ( $\beta_3 \times 100$ )	0.233 (0.604)	-0.012 (0.011)
Constant ( $\alpha$ )	-201.61 (113.34)	3.832 (2.215)
Estimated peak age	25.59	21.15
95% confidence interval	[24.78, 26.40]	[17.23, 25.08]
Player fixed effects	Yes	Yes
N of players	1,152	715
N of player-seasons	4,703	2,481
Within $R^2$	0.038	0.051

Notes.- Estimates of Equation (2): minutes per game concern 2007-18; [WhoScored.com](#) ratings concern 2013-19. We could statistically reject the cubic specification in favour of a quadratic for these models, but for consistency throughout the paper and across models we still prefer to report the cubic estimates (see Figures 5B and 5D). Standard errors in parentheses, as well as the 95% confidence interval of the age peaks, are calculated using bootstrapping with 100 repetitions.

## 4. Age-wage profiles

### 4.1 Empirical strategy

Since we have salary data for practically all workers playing in MLS over thirteen years, we can robustly estimate average age-wage profiles. We first estimate age-wage profiles using a worker-level fixed effects model to account for selection effects and relevant unobservable worker-firm-specific factors, as in Section 3 & Equation (2):

$$w_{i,t} = \alpha + \beta \text{Age}_{i,t} + \lambda_i + \varepsilon_{i,t} \quad (\text{Age-wage profile}), \quad (5)$$

where  $w_{i,t}$  gives the log wage of worker  $i$  in year  $t$ . We estimate Equation (5) for two separate samples: one including designated players (DPs) and one without.

We also investigate how worker age-wage and age-productivity profiles differ, by estimating an *age-wage premium* profile, adding extra individual performance variables as regressors to the model given by Equation (5). This can account for the part of a worker's salary variance over time that is explained by their on-pitch performance. Since wages are determined at the beginning of an MLS season, we use lagged performance variables, including the main productivity measure from Section 3: minutes played per game; as well as goals scored, shots on/off goal, assists, fouls committed/conceded and yellow/red cards for outfield players, and saves for goalkeepers, where all these variables are normalised per 90 minutes played, and are included as cubic polynomials to fit their relationship with wages. We thus estimate the following:

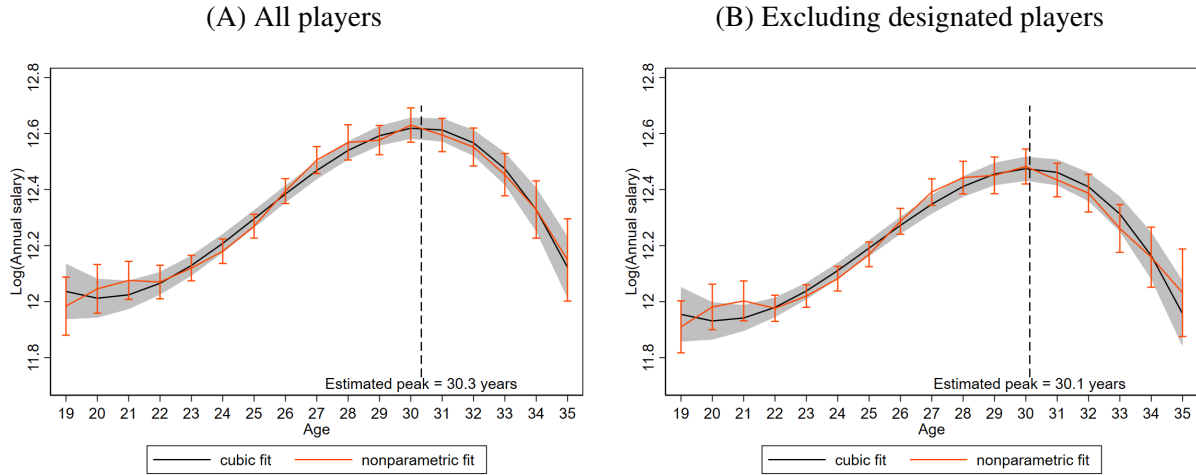
$$w_{i,t} = \alpha + \beta \text{Age}_{i,t} + \delta \mathbf{x}_{i,t-1} + \lambda_i + \varepsilon_{i,t} \quad (\text{Age-wage-premium profile}), \quad (6)$$

where  $\mathbf{x}_{i,t}$  denotes a vector containing the set of worker- and time-specific performance variables described above. We estimate Equation (6) separately for outfield players and goalkeepers (for detailed evidence on the salary determination of professional goalkeepers and how this differs from outfield players, see [Berri et al., 2023](#)).

### 4.2 Results

Figure 6 and Table 3 display the results from estimating Equation (5) using least squares. Wages are estimated to peak at age 30, which is significantly higher than the estimated productivity peaks shown in Table 2. Workers thereafter tend to experience a sharp drop-off in their wages, such that by age 35 wages have fallen by 50 log points compared with the peak and are almost back to the same level as at age 19 (Figure 6A). Excluding designated players from the estimation sample does not alter these results significantly (Figure 6B), such that the late peak age of wages compared with productivity does not appear to be driven by the designated players in MLS.

FIGURE 6: Estimated *age-wage* profiles using worker-level fixed effects, 2007-2019



Notes.- Estimates of Equation (5). Wages in “MLS 2019 prices” (see Section 2). Shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

TABLE 3: Estimated *age-wage* profiles using worker-level fixed effects, 2007-2019

	All workers (I)	Excl. designated players (II)
Age ( $\beta_1$ )	-2.103 (0.251)	-2.002 (0.248)
Age <sup>2</sup> ( $\beta_2 \times 10$ )	0.869 (0.094)	0.829 (0.094)
Age <sup>3</sup> ( $\beta_3 \times 10$ )	-0.115 (0.012)	-0.110 (0.02)
Constant ( $\alpha$ )	28.50 (2.206)	27.61 (2.179)
Estimated peak age	30.34	30.13
95% confidence interval	[30.00, 30.67]	[29.80, 30.45]
Player fixed effects	Yes	Yes
<i>N</i> of worker	1,283	1,242
<i>N</i> of worker-years	5,245	4,925
Within <i>R</i> <sup>2</sup>	0.139	0.128

Notes.- Estimates of Equation (5), with the dependent variable of log wages in “MLS 2019 prices” (see Section 2). We can statistically reject a quadratic specification in favour of a cubic for these models (see Figure 6). Standard errors in parentheses and the 95% confidence interval of the age peaks calculated using bootstrapping with 100 repetitions.

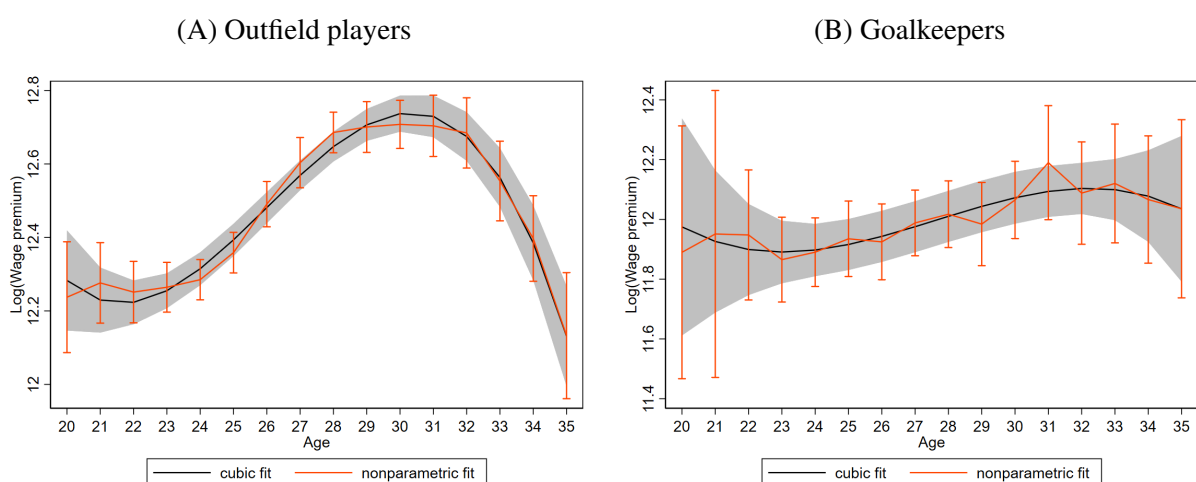
The timings of the peak and drop-off in wages differ substantially from those for productivity shown in Section 3. As Figure 7A shows, when we also control for observable productivity variables by estimating Equation (6), we find that the *age-wage premium* increases throughout

most of an MLS outfield player's career, until it reaches a maximum in their early 30s, and then falls sharply up to age 35. We only observe this pattern for outfield players, finding no significant age-wage premium among the relatively small number of goalkeepers in the League (Figure 7B). The life-cycle profile of the outfield age-wage premium is on the surface a puzzle, since we might expect firms to compete for a player by offering higher wages up until the point where their salary reflects their productivity, implying that wages and on-pitch productivity would peak at the same age and the life-cycle profiles then run parallel to one another. In the next section, we discuss and evaluate some possible solutions to this puzzle.

As discussed in Section 2, we should question how our findings from the MLS might generalise to other high-wage football leagues. Therefore, as a robustness check, we repeat our analysis using data from the German Bundesliga, which as of the 2022/23 is the third highest ranked football league in Europe (see Online Appendix B).<sup>10</sup> Although the Bundesliga salary data are incomplete and unofficial, we find similar patterns to those from MLS, providing some limited reassurance of external validity.

We also perform two additional exercises to test the robustness of the estimated age-wage premium profiles. First, to test whether the shape of the profile is driven by players with only short stints in MLS, we repeat the estimation of Equation (6) using only outfield players who are observed for a minimum of five seasons (52% (367/707) of the players used to estimate Figure 7A). The results are generally consistent with the larger original sample of players (Online Appendix Figure A5A). Second, to test whether the puzzle is driven only by within-club contracting, we estimate Equation (6) on outfield players who spent their entire MLS career in one club (42% (295/707) of the players used to estimate Figure 7A). The results are also generally consistent with those using the original sample (Online Appendix Figure A5B).

FIGURE 7: Estimated *age-wage premium* profiles using worker-level fixed effects, 2008-2019



Notes.- Estimates of Equation (6), for the 707 outfield players and 123 goalkeepers who are observed at least twice with two consecutive seasons in MLS (e.g., a continuous spell of 3 seasons). Wages in “MLS 2019 prices” (see Section 2). Shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

<sup>10</sup>See UEFA coefficient rankings: <https://www.uefa.com/nationalassociations/uefarankings/country>.

## 5. The age-wage-productivity puzzle

Our findings so far provide robust evidence that an age-wage premium exists in MLS: although workers in this labour market peak in productivity during the middle of their careers, their wages continue to increase for several years after, before then dropping sharply. In this section, we discuss four potential explanations for this discrepancy between the estimated age-productivity and age-wage profiles. The first explanation may be specific to this particular type of labour market: (1) regulations in MLS mandating annual salary increases for lower-paid players (who are likely to be younger).

We then consider three more general explanations as to why younger workers may be ‘underpaid’: (2) “superstar” effects, whereby mid-career players are paid more because they are more popular and attract greater audiences (despite not being more talented); (3) the existence of some aspect of productivity or human capital that is not reflected in our measures of productivity, but is related to age (leadership qualities, for example); and (4) a “talent discovery” effect, whereby older workers earn more because their level of talent is better known. We also explore a theory for why *older* workers (near age 35) appear to be underpaid relative to ability. We consider whether greater injury-proneness at later stages in a player’s career increases the variance of their performance, which may cause risk-averse clubs to underpay for older talent.

Before turning to these four explanations, it is worth commenting on another mechanism that has been suggested as a driver of age-wage profiles in other markets, which we argue should not obviously apply to the MLS setting. This theory, initially suggested in a series of papers by Lazear (1979, 1981), suggests that long-term contracts between firms and workers may entail *delayed compensation* to avert shirking and malfeasance. This has been used to explain why workers may be paid less than their marginal contribution to an employer when young, and vice versa when old. However, for Lazear’s theory to explain an age-wage-productivity puzzle two conditions must be met. First, for delayed compensation to be a non-shirking device, it is important that the effort and performance of workers is difficult to monitor – a condition which trivially does not apply for professional footballers. Second, the theory requires that the relationships formed between firms and workers are expected by both parties to be long-term. This is not the case in professional football, where the generality and visibility of player performance leads to high turnover and a competitive market for talent, such that very long-term relationships between clubs and players are unlikely from the outset. In a study of German Bundesliga footballers, Buraimo et al. (2015) found evidence suggesting that the pattern of selection on to long-term contracts was explained by performance and employer risk aversion rather than moral hazard concerns, which points away from Lazear’s theory.

### 5.1 Regulation

MLS is a single-entity that owns a stake in all teams in the League. Players sign a contract with MLS, which specifies an initial guaranteed period, during which the contract cannot be terminated

due to poor performances or injury. This is followed by up to three option years, when MLS has the right to extend the contract. There are a number of clauses in the collective agreement between MLS and the players' association concerning players' salaries. In particular, regulation surrounding salary increases are potentially relevant to the age-wage premium. The 2015-19 CBA specified that players earning below \$150,000 must receive an increase of 5% per year in each year of their contracts. Between 2005 and 2009, this clause applied to all players earning less than \$60,000.<sup>11</sup> This means that, for some players at least, salaries increase every year, regardless of their productivity.

In addition, the CBA specifies minimum salaries for each season. The minimum salary for a player aged under 24 on a team's reserve roster is lower than for an older player on the main roster. For example in 2018, the minimum salary was \$54,500 for players under 24 and \$67,500 for older players. However, most players earn over the minimum. For example, all but 26 of 537 players older than 24 were paid more than \$67,500 for the 2018 season. Most players also received annual salary increases that were far greater than the minimum; the median salary increase was 9.5% and the annual nominal change in salary was greater than 5% for 66% of player-year observations in our estimation samples. Therefore, we would not expect the age-wage premium profile to be a result of salary regulation. If this were the case, then MLS would not hire older players, preferring cheaper younger and very old players.

However, to confirm whether the age-wage premium is determined by salary regulation, we perform a further robustness check. We repeat our estimation of the age-wage profile, including a dummy variable for each period covered by the three CBAs, i.e., one dummy variable for the years 2007 to 2009, one for the years 2010 to 2014, and one for the years 2015 to 2019. We also include the interaction between a player's age and the CBA period, such that we estimate the following model:

$$w_{i,t} = \alpha + \beta \text{Age}_{i,t} + \eta \text{CBA}_t + (\phi \text{Age}_{i,t})' \text{CBA}_t + \lambda_i + \varepsilon_{i,t} , \quad (7)$$

where  $\beta$  and  $\text{Age}$  are as in the non-parametric version of Equation (2) (with a dummy for each age group).  $\text{CBA}_t$  is a column vector of dummy variables for each CBA period and  $\eta$  is a column vector of coefficients.  $\phi$  is a  $3 \times 16$  matrix of age and CBA period specific interaction coefficients. If regulation governed by the different CBAs had an effect on the age-wage profile, then we would expect that the interaction coefficients in  $\phi$  are significantly different from zero. This is not the case.<sup>12</sup> In other words, there is no evidence that changes in regulation explain the age-wage profile and the differences between the age-productivity and age-wage profiles.

<sup>11</sup>MLSPA and MLS signed a Memorandum of Understanding (which is not publicly available) just before the 2010 season began, but never wrote a formal CBA governing the 2010 to 2014 seasons.

<sup>12</sup>Note that, since we always study wages in 2019 MLS prices, we would not expect the  $\eta$  coefficients to be significant, which is the case. Full estimation results are available upon request.

## 5.2 Superstar effects

Substantial evidence exists that top earners in industries such as media and sports may be paid for factors other than pure talent (Carrieri et al., 2018; Hoffman and Opitz, 2019; Filimon et al., 2011). For example, since a football team’s revenue ultimately comes from fans attending and viewing matches, we might expect other factors, such as the charisma and popularity of its players, to affect the profitability of a team and thus the marginal revenue product of individual players. There is evidence that ‘superstars’ affect fans’ willingness to pay for tickets to sporting events (e.g., Hausman and Leonard, 1997; Kaplan, 2022), including specifically in MLS (Coates et al., 2016; Jewell, 2017), and the wages of football players (Carrieri et al., 2018; Scarfe et al., 2021). It is possible, therefore, that superstar effects are biasing our age-productivity estimates, insofar as the level of superstardom could be correlated with age in a way that cannot be explained by on-pitch performances. Indeed, such an age effect is plausible if part of a player’s popularity comes from building up a reputation, or a fan base, over time.

To test whether superstar effects are a possible explanation of our puzzle, we use Wikipedia page views as a proxy for player popularity. We collected data on the annual Wikipedia page views of MLS players’ English language profiles in the years 2016-2018, using the Pageview Application Programming Interface (API), a tool for querying the Wikipedia Foundation page views data. Online Appendix Figure A6 shows the distribution of our proxy popularity measure over player-seasons, by plotting the density of the logarithm of annual page views. To analyse the effect of age on the popularity of a player, we regress the logarithm of their profile page views on age using non-parametric and parametric specifications, as in the previous sections. Online Appendix Figure A7A shows the results for the ‘OLS’ specification, which does not control for worker-level fixed effects. There is a steep age-popularity gradient, with older workers attracting significantly more page views than younger workers. However, it appears that this result is mainly driven by selection, rather than the effect of age on popularity. When controlling for worker-level fixed effects, in Online Appendix Figure A7B, we find that page views track players’ performance over a career more closely than their wages. The page views tend to peak relatively early in a career and fade with age.<sup>13</sup> Taken together, these results suggest that a superstar effect may have a role in explaining why older workers earn more before accounting for selection effects, which is unsurprising given MLS’s strategy of purchasing ageing superstars. However, it does not appear that popularity explains our finding that the age-wage premium is increasing for a given worker as they age, as our proxy measure of popularity tends to track individual performance more closely than wages over a career in MLS.

---

<sup>13</sup>As a robustness check, we also estimate the model after dropping designated players from the sample, in order to make sure that our results are not driven by the first-year buzz of newly signed designated players, which has been shown to fade with time (Jewell, 2017). However, this does not alter the results in a meaningful way – results are available on request.



### 5.3 Unobserved productivity or human capital

It is possible that MLS players accumulate human capital during their careers that is not reflected in our data. Our ‘direct’ productivity measures should do a good job at picking up on-the-pitch performance; the [WhoScored.com](#) ratings are based on many measurable performance indicators, such as completed passes, goals scored, interceptions etc., and minutes played should also account for general fitness and on-the-pitch leadership qualities that the manager observes but the [WhoScored.com](#) ratings fail to pick up. However, it is still possible that workers in their 30s contribute more to the performance of the firm with *off-the-pitch* qualities, such as responsibilities during training or in the dressing room. This could explain why they seem to be overpaid relative to their ability.

To test whether players around age 30 – who we have labelled as ‘overpaid’ – contribute positively to team performance, we use the variation in team-specific age profiles (see Figure 1). Using the estimated age-coefficients from the non-parametric version of Equation (6) (Figure 7) and the sample of outfield players, we construct a team×age specific variable  $WP$ , which reflects the age-induced wage premium of a firm, i.e., how much it appears to ‘overpay’ its workers based on the age-structure of its roster alone:

$$\widehat{WP}_{j,t} = \sum_{i=1}^{N_{j,t}} \hat{\beta}_{A(i,j,t)} / N_{j,t} , \quad (8)$$

where subscripts  $i$ ,  $j$  and  $t$  correspond to worker, firm and year, respectively, and  $\text{Age} = A(i, j, t)$  is an indicator function for the age of worker  $i$  at firm  $j$  in year  $t$ , that relates to a coefficient estimated from the age-wage premium model.  $N_{j,t}$  denotes the roster size (excluding goalkeepers). We then regress average points per game over a season,  $ppg_{j,t}$ , on  $\widehat{WP}_{j,t}$ .<sup>14</sup> We also control for the firm’s average salary,  $\text{WageBill}_{j,t}$ , which is the log of the team’s average wage (in “2019 MLS dollars”, see Section 2) in a given season.<sup>15</sup> Controlling for the firm-specific wages is important, since we find that the most productive workers in the League *all-in-all*, before controlling for selection, are indeed in their 30s (see Figure 5A). Hence, having an age structure with more young (or old) workers may only be a cheaper way for a team and MLS as a whole to get higher average quality performance for a given budget. We estimate the following model:

$$ppg_{j,t} = \alpha + \gamma \widehat{WP}_{j,t} + \phi \text{WageBill}_{j,t} + \varepsilon_{j,t} , \quad (9)$$

with results reported in Online Appendix Table A1. Qualitatively, these estimates are in line with the findings throughout our study; having an age structure with a higher wage-premium is associated with winning fewer football matches, after controlling for the salary bill. This again poses the puzzle of why firms do not focus their recruitment strategies on relatively young or old workers. However, this result must be interpreted with caution since the estimated coefficients are

<sup>14</sup>Points per game are calculated as per MLS’s table, with three points for a win and one point for a draw.

<sup>15</sup>We find the average salary a more robust measure than aggregate salary, since we are missing a small number of player-season observations

not significantly different from zero. Although there is a reasonable amount of variance over firms and years in both points per game and the age-induced wage premiums in the data, as shown in Online Appendix Figure A8, which displays kernel density estimates for both of these variables, the exercise has relatively low statistical power, with only 217 firm-year observations. To put the magnitudes of the estimates in perspective, the estimated  $\hat{\gamma} = -0.29$  suggests that the most underpaid team imaginable (consisting only of 35-year-olds) would on average earn 0.175 more points per game than the most overpaid team imaginable (consisting only of 30-year-olds).

## 5.4 Talent discovery and risk aversion

A fourth potential explanation for our puzzle is that firms are *risk-averse*, which causes older workers to earn a premium if they are considered ‘safer bets’ than younger ones. If a worker’s productivity over time has a large component of idiosyncratic variance, then it is plausible that firms will pay more to workers with more available past performance data. There is evidence from sports labour markets that firms pay more to consistent performers (e.g., [Deutscher et al., 2017](#); [Özdemir et al., 2022](#)), at the same time as over-valuing recent performance (e.g., [Healy, 2008](#)). In other labour markets, [Kuhnen and Oyer \(2016\)](#) found that firms are more likely to hire MBA graduates who have previously worked in the same industry, suggesting some level of risk aversion.

This concept is similar to a mechanism proposed by [Terviö \(2009\)](#). In his model, talent is only revealed after starting work. [Terviö](#) shows that, when workers are unable to commit to long-term wage contracts, firms prefer to hire experienced workers whose talent is known, even if inexperienced workers may be more talented on average.<sup>16</sup> This theory has proved difficult to test empirically. An exception is [Peeters et al. \(2022\)](#), who studied the hiring of football managers in England, finding that employers seem to prefer lower ability managers with experience over novice managers with higher expected ability.

Our data provide an opportunity to test this ‘talent discovery’ hypothesis in a high-wage labour market. Talent in this market is revealed as workers age, through more seasons of past performance being easily visible and recorded. For younger workers, however, an especially strong season provides less proof of high permanent ability, and a player could be a ‘one-season wonder’ (see [DeSchrive \(2007\)](#) for analysis of an infamous case in MLS). We test whether observing more past performance data among younger players helps to predict future productivity, by regressing our productivity measures on their lagged values: if idiosyncratic variance is important, adding more past seasons to this quasi-forecasting model should add significantly more predictive power. Limiting our estimation sample to workers aged 30 or younger who have had at least a four-season long stint in MLS with non-missing performance observations, giving 412 unique workers for minutes played per game and 196 workers for [WhoScored.com](#) ratings, Table 4 displays the

---

<sup>16</sup>In a sense, this is the opposite of [Lazear \(1998\)](#)’s theory, which suggests that firms value riskier workers more, since those who perform better than their expected productivity can be retained, and those who perform worse can be fired. That theory has found some empirical support. For example, [Bollinger and Hotchkiss \(2003\)](#) find that baseball players with more variable performance are paid more.

results of estimating these simple forecasting regression models. For minutes played, although lags from more than one season ago are significant, they are worse predictors of performance than the most recent season, and adding more lags does not increase the R-squared or decrease the root mean squared error of the regression much at all. One interpretation of these results is that the persistent component in a worker's performance process, when measured by minutes played per game, is much larger than the idiosyncratic component, such that performance data from more than one season ago do not help to predict a worker's future performance. Perhaps this result is not too surprising, since the season-level data we have access to is already aggregated over dozens of matches. It is, however, notable that the estimated constant of the regression model attenuates substantially as lags of minutes per game are added. This implies that the efficiency of the forecasting model, or its bias towards under-predicting high productivity players, improves as more seasons are used in a weighted average to predict future minutes per game.

TABLE 4: Estimation results from autoregressive models of minutes played per game and [WhoScored](#) ratings

	Minutes per game			<a href="#">WhoScored.com</a> rating		
	(I)	(II)	(III)	(IV)	(V)	(VI)
First lag	0.586 (0.024)	0.521 (0.031)	0.516 (0.031)	0.375 (0.046)	0.300 (0.055)	0.262 (0.059)
Second lag		0.109 (0.033)	0.049 (0.038)		0.229 (0.077)	0.161 (0.075)
Third lag			0.105 (0.030)			0.183 (0.055)
Constant	15.43 (1.322)	13.44 (1.446)	12.07 (1.444)	4.205 (0.313)	3.155 (0.456)	2.628 (0.516)
<i>N</i> of players	412	412	412	196	196	196
<i>N</i> of player-seasons	1,134	1,134	1,134	407	407	407
$R^2$	0.324	0.332	0.339	0.125	0.158	0.185
Root MSE	22.48	22.25	22.01	0.286	0.280	0.276

Notes.- minutes per game concern 2007-18; [WhoScored.com](#) ratings concern 2013-19. Sample includes players who played at least consecutive 4 seasons in MLS when aged 15-30. Robust standard errors in parentheses.

Interestingly, using [WhoScored.com](#) ratings as a dependent variable paints a slightly different picture. In this case, lagged performance effects are significant and large up until and including the third lag, and including three lags increases the R-squared by around fifty percent compared to just including one lag. This suggests that performance on-the pitch, as measured by [WhoScored.com](#) ratings, has a significant year-to-year idiosyncratic component. The constant of the regression model for [WhoScored.com](#) ratings attenuates as lags are added, even more so than for minutes per game, again implying that more past data allows for less biased forecasts of future performances. In this case, firms may be hesitant to sign young players as they cannot predict accurately or

efficiently how good they are from the limited data available to them. We think this provides the most plausible explanation to the age-wage-productivity puzzle that we have considered for the apparently low wages received by young players, although the evidence is far from conclusive.

While the model estimates in Table 4 provide a basic test of the hypothesis that forecasting the ‘true’ ability of a worker requires several years worth of data, there is an alternative mechanism through which risk aversion in particular could explain the inverted u-shape of the wage premium: if mid-career workers are *on average* more consistent performers than younger and older ones, then firms may be willing to pay mid-career workers a risk premium. This is not implausible, since younger players may be more variable due to a lack of experience, whereas older players may be more prone to injuries (or, more generally, to declines in health that affect their productivity). We test this hypothesis by directly estimating the year-specific variance of performance as a function of age, using a two-step procedure. First, we estimate age-productivity profiles, including worker-level fixed effects as in Section 3, and collect the residuals. Second, we regress those squared residuals on age, again including worker-level fixed effects to account for the possibility that some workers are inherently more consistent performers than others. Online Appendix Figure A9 displays the results from those second-step regressions. We do not find clear evidence for an inverted u-shaped age profile in the variability of performance. For minutes played per game (Online Appendix Figure A9A), the profile appears to be increasing but at a diminishing rate throughout workers’ careers. For [WhoScored.com](#) ratings (Online Appendix Figure A9B), there is no evidence of an age profile in the variability of performance over seasons. Hence, there is no evidence that mid-career workers being on average more consistent can explain the age-wage premium puzzle.

## 6. Conclusion

There are a number of reasons why it is difficult for researchers to estimate robustly the effects that a worker’s age has on their wages and productivity. In this paper, we used data from Major League Soccer in the United States to estimate professional football players’ age-productivity and age-wage profiles. Our data provide us with both observable productivity and salary measures, a combination that is not often available to researchers. We also observe workers’ performance and salary over their whole career in MLS, allowing us to estimate age-productivity and age-wage profiles that control for unobserved heterogeneity and selection in and out of the market. We found that age-productivity profiles peak significantly earlier than age-wage profiles.

A hump-shaped age-wage profile is well established in many labour markets ([Mincer, 1958](#); [Huggett et al., 2011](#); [Rupert and Zanella, 2015](#)). In terms of productivity, our findings correspond with those of [Bertoni et al. \(2015\)](#), [Castellucci et al. \(2011\)](#), and [Hakes and Turner \(2011\)](#), from other professional sports. They also found that peak productivity would appear to happen significantly later when looking at a cross-section of workers and failing to account for selection effects. Studies in other non-sports markets have also found that the productivity of a worker

initially increases with experience, before either flattening off or declining. The timing of the peak and the extent of the decline, as one would expect, appear to depend on the setting ([Oster and Hamermesh, 1998](#); [Börsch-Supan et al., 2021](#)).

The structure of our dataset allows us to estimate robustly an *age-wage premium* profile, that also addresses selection effects and unobserved worker-specific factors that are valued by employers but unrelated to age. Although data limitations in most labour markets mean that the accurate estimation of an age-wage premium is difficult, the underpayment of younger workers has been observed in other (primarily manufacturing) industries and settings ([Dostie, 2011](#); [Ilmakunnas and Maliranta, 2005](#); [Cardoso et al., 2011](#)). Our findings provide new evidence for such a premium in a high-wage labour market that has limited expectations of long-term employment relationships.

We investigated a number of plausible reasons for the estimated differences between the age profiles of productivity and wages in MLS: the role of regulation or wage bargaining that are specific to this labour market; the possibility that there are unobserved elements of productivity correlated with age; and uncertainty causing firms to prefer players who are past their peak but still reasonably far from the end of their careers. We do not find convincing evidence that any of these can explain the shape of the estimated age-wage premium profile. However, we do find that more available data on past productivity is useful in predicting future productivity among young players. This suggests that the productivity of older workers may be better known, making them a better investment for firms that are either risk averse or credit constrained. Our findings provide some further support for a theoretical model proposed by [Terviö \(2009\)](#), in which firms respond to uncertainty in the ability of younger workers by excessively bidding for older workers whose talent is better known. Empirical evidence on the effect of information about past performance on hiring decisions is currently limited (exceptions include [Peeters et al., 2022](#); [Pallais, 2014](#); [Stanton and Thomas, 2016](#)). This is a promising avenue for future research into the age-wage premiums that appear to exist across a range of different labour markets. Unfortunately, we do not have reliable, representative data on the length of contracts in MLS or professional football more generally, although they are widely known to be rarely longer than a few years and commonly even shorter. But there is some relevant evidence from Major League Baseball that employers tend to overpay more for talent when offering longer contracts ([Solow and Krautmann, 2020](#); [Walters et al., 2017](#)). If MLS teams offer older or younger workers shorter contracts compared with mid-career workers, then this would align with a notion that firms overpay and perhaps over-compete when recruiting or retaining the latter because they are less risky investments - insuring against the challenges of replacing key units of known talent in a business that tends to value short-term success. In future, good data on contract lengths will perhaps allow researchers to test this mechanism.

A perspective we have not considered is whether individual worker performances are endogenous to the distribution of salaries in the firm or market. In MLS, for example, [Coates et al. \(2016\)](#) demonstrated that greater within-team salary dispersion is associated with worse team performance in the League, aligning with a theory of cohesive firms ([Levine, 1991](#)), and matching

results from other sports leagues (e.g., [Depken II, 2000](#); [Jane, 2010](#)). Further, the presence of pay transparency should allow the players and other agents to form views on who is over or underpaid, thus changing behaviour accordingly. [Flynn \(2022\)](#) has shown that the introduction of pay transparency in the National Hockey League led to apparently underpaid players switching toward more offensive play because this was more valued by employers. It is also well-known that offensive attributes are highly valued in professional football (e.g., [Lucifora and Simmons, 2003](#)). We controlled for these concerns to some extent in the estimation of our age-wage premium profiles, but future research could ask, for instance, whether experience and tenure in MLS enables mid-career players to take up more offensive roles, before their legs give out and they return to more defensive roles toward the end of their careers. Again, a benefit of using sports data to consider these questions is the availability of data about player behaviour, such as their choice of style and the specific roles they fulfil within teams.

## References

- Adler, M.** 1985. "Stardom and Talent." *American Economic Review*, 75(1): 208–212.
- Amand, M., A. Chéron, F. Pelgrin, and A. Terriau.** 2023. "Soccer labour market equilibrium and efficient training of talents." *European Economic Review*.
- Barth, M., A. Güllich, C. A. Forstinger, T. Schlesinger, F. Schröder, and E. Emrich.** 2021. "Retirement of professional soccer players – A systematic review from social sciences perspectives." *Journal of Sports Sciences*, 39(8): 903–914.
- Berri, D., D. Butler, G. Rossi, R. Simmons, and C. Tordoff.** 2023. "Salary determination in professional football: empirical evidence from goalkeepers." *European Sport Management Quarterly*.
- Bertoni, M., G. Brunello, and L. Rocco.** 2015. "Selection and the age – productivity profile: Evidence from chess players." *Journal of Economic Behavior & Organization*, 110 154–156.
- Blass, A. A.** 1992. "Does the Baseball Labor Market Contradict the Human Capital Model of Investment?" *The Review of Economics and Statistics*, 74(2): 261–268.
- Bollinger, C., and J. Hotchkiss.** 2003. "The Upside Potential of Hiring Risky Workers: Evidence from the Baseball Industry." *Journal of Labor Economics*, 21(4): 923–944.
- Börsch-Supan, A., C. Hunkler, and M. Weiss.** 2021. "Big data at work: Age and labor productivity in the service sector." *The Journal of the Economics of Ageing*, 19.
- Bradbury, J. C.** 2019. "Determinants Of Revenue In Sports Leagues: An Empirical Assessment." *Economic Inquiry*, 57(1): 121–140.
- Bradbury, J. C.** 2021. "Financial Returns in Major League Soccer." *Journal of Sports Economics*, 22(8): 921–945.
- Brown, J. N.** 1989. "Why Do Wages Increase with Tenure? On-the-Job Training and Life-Cycle Wage Growth Observed within Firms." *American Economic Review*, 79(5): 971–91.
- Buraimo, B., B. Frick, M. Hickfang, and R. Simmons.** 2015. "The Economics of Long-term Contracts in the Footballers' Labour Market." *Scottish Journal of Political Economy*, 62(1): 8–24.
- Bureau of Labor Statistics.** 2021. "Number of Jobs Held, Labor Market Activity, and Earnings Growth Over Two Decades: Results from a Longitudinal Survey Summary.." News Release.
- Burguet, R., and J. Sákovics.** 2019. "Bidding For Talent In Sport." *Economic Inquiry*, 57(1): 85–102.
- Butler, D., and D. Coates.** 2022. "Position Premium in Major League Soccer." *International Journal of Sport Finance*, 17(4): .
- Card, D., A. R. Cardoso, J. Heining, and P. Kline.** 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36(S1): 13–70.
- Cardoso, A. R., P. Guimarães, and J. Varejão.** 2011. "Are Older Workers Worthy of Their Pay? An Empirical Investigation of Age-Productivity and Age-Wage Nexuses." *De Economist*, 159(2): 95–111.



- Carrieri, V., F. Principe, and M. Raitano.** 2018. “What makes you ‘super-rich’? New evidence from an analysis of football players’ wages.” *Oxford Economic Papers*, 70(4): 950–973.
- Castellucci, F., M. Padula, and G. Pica.** 2011. “The age-productivity gradient: Evidence from a sample of F1 drivers.” *Labour Economics*, 18(4): 464–473.
- Coates, D., B. Frick, and T. Jewell.** 2016. “Superstar Salaries and Soccer Success: The Impact of Designated Players in Major League Soccer.” *Journal of Sports Economics*, 17(7): 716–735.
- De Hek, P., and D. van Vuuren.** 2011. “Are older workers overpaid? A literature review.” *International Tax and Public Finance*, 18(4): 436–460.
- Depken II, C. A.** 2000. “Wage disparity and team productivity: evidence from major league baseball.” *Economics Letters*, 67(1): 87–92.
- DeSchrive, T. D.** 2007. “Much Adieu About Freddy: Freddy Adu and Attendance in Major League Soccer.” *Journal of Sport Management*, 21(3): 438–451.
- Deutscher, C., O. Gürtler, J. Prinz, and D. Weimar.** 2017. “The payoff to consistency in performance.” *Economic Inquiry*, 55(2): 1091–1103.
- Dostie, B.** 2011. “Wages, Productivity and Aging.” *De Economist*, 159 139–158.
- Dustmann, C., and S. C. Pereira.** 2008. “Wage growth and job mobility in the united kingdom and germany.” *ILR Review*, 61(3): 374–393.
- Ericson, T.** 2000. “The Bosman Case.” *Journal of Sports Economics*, 1(3): 203–218.
- Fair, R. C.** 2008. “Estimated Age Effects in Baseball.” *Journal of Quantitative Analysis in Sports*, 4(1): .
- Fair, R. C., and E. H. Kaplan.** 2018. “Estimating Aging Effects in Running Events.” *The Review of Economics and Statistics*, 100(4): 704–711.
- Ferguson, D. G., J. C. H. Jones, and K. G. Stewart.** 2000. “Competition Within A Cartel: League Conduct And Team Conduct In The Market For Baseball Player Services.” *The Review of Economics and Statistics*, 82(3): 422–430.
- Filimon, N., J. López-Sintas, and C. Padrós-Reig.** 2011. “A test of Rosen’s and Adler’s theories of superstars.” *Journal of Cultural Economics*, 35 137–161.
- Flynn, J.** 2022. “Salary disclosure and individual effort: Evidence from the National Hockey League.” *Journal of Economic Behavior & Organization*, 202(C): 471–497.
- Fort, R.** 2000. “European and North American Sports Differences(?)” *Scottish Journal of Political Economy*, 47(4): 431–455.
- Frick, B.** 2007a. “The Football Players’ Labor Market: Empirical Evidence from the Major European Leagues.” *Scottish Journal of Political Economy*, 54(3): 422–446.
- Frick, B.** 2007b. “The football players’ labor market: Empirical evidence from the major european leagues.” *Scottish Journal of Political Economy*, 54(3): 422–446.
- Frick, B.** 2009. “Globalization and Factor Mobility: The Impact of the “Bosman-Ruling” on Player Migration in Professional Soccer.” *Journal of Sports Economics*, 10(1): 88–106.

- Goddard, J., P. J. Sloane, and J. O. S. Wilson.** 2012. “The Bosman Ruling and Labor Mobility in Football (Soccer).” In *The Oxford Handbook of Sports Economics: The Economics of Sports Volume 1*. Oxford University Press.
- Hakes, J., and C. Turner.** 2011. “Pay, productivity and aging in Major League Baseball.” *Journal of Productivity Analysis*, 35(1): 61–74.
- Hanoch, G., and M. Honig.** 1985. ““True” Age Profiles of Earnings: Adjusting for Censoring and for Period and Cohort Effects.” *The Review of Economics and Statistics*, 67(3): 383–394.
- Hausman, J. A., and G. K. Leonard.** 1997. “Superstars in the National Basketball Association: Economic Value and Policy.” *Journal of Labor Economics*, 15(4): 586–624.
- Healy, A.** 2008. “Do Firms Have Short Memories?: Evidence From Major League Baseball.” *Journal of Sports Economics*, 9(4): 407–424.
- Hellerstein, J. K., D. Neumark, and K. R. Troske.** 1999. “Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations.” *Journal of Labor Economics*, 17(3): 409–446.
- Hoffman, K., and C. Opitz.** 2019. “Talent and publicity as determinants of superstar incomes: empirical evidence from the motion picture industry.” *Applied Economics*, 51(13): 1383–1395.
- Holmes, B., and I. G. McHale.** 2023. “Forecasting football match results using a player rating based model.” *International Journal of Forecasting*.
- Huggett, M., G. Ventura, and A. Yaron.** 2011. “Sources of Lifetime Inequality.” *American Economic Review*, 101(7): 2923–54.
- Ilmakunnas, P., and M. Maliranta.** 2005. “Technology, Labour Characteristics and Wage-productivity Gaps.” *Oxford Bulletin of Economics and Statistics*, 67(5): 623–645.
- Jane, W.-J.** 2010. “Raising salary or redistributing it: A panel analysis of Major League Baseball.” *Economics Letters*, 107(2): 297–299.
- Jewell, R. T.** 2017. “The Effect of Marquee Players on Sports Demand: The Case of U.S. Major League Soccer.” *Journal of Sports Economics*, 18(3): 239–252.
- Kahn, L. M.** 2000. “The Sports Business as a Labor Market Laboratory.” *Journal of Economic Perspectives*, 14(3): 75–94.
- Kaplan, S. M.** 2022. “Putting a price on popularity: Evidence from superstars in the National Basketball Association.” *Economic Inquiry*, 60(3): 1357–1381.
- Késenne, S.** 2006. “The Win Maximization Model Reconsidered: Flexible Talent Supply and Efficiency Wages.” *Journal of Sports Economics*, 7(4): 416 – 427.
- Kesenne, S.** 2015. “The single entity status of a sports league.” *Journal of Sports Economics*, 16(8): 811–818.
- Kuethe, T., and M. Motamed.** 2010. “Returns to Stardom: Evidence From U.S. Major League Soccer.” *Journal of Sports Economics*, 11(5): 567–579.
- Kuhnen, C. M., and P. Oyer.** 2016. “Exploration for Human Capital: Evidence from the MBA Labor Market.” *Journal of Labor Economics*, 34(S2): S255–S286.

- Lazear, E.** 1998. "Hiring risky workers." In *Internal Labour Market, Incentives, and Employment*. Eds. by S. Ohashi, and T. Tachibanaki, New York: St. Martin's.
- Lazear, E. P.** 1979. "Why Is There Mandatory Retirement?" *Journal of Political Economy*, 87(6): 1261–1284.
- Lazear, E. P.** 1981. "Agency, Earnings Profiles, Productivity, and Hours Restrictions." *The American Economic Review*, 71(4): 606–620.
- Levine, D. I.** 1991. "Cohesiveness, productivity, and wage dispersion." *Journal of Economic Behavior & Organization*, 15(2): 237–255.
- Lucifora, C., and R. Simmons.** 2003. "Superstar Effects in Sport: Evidence From Italian Soccer." *Journal of Sports Economics*, 4(1): 35–55.
- Major League Soccer.** 2020. "Media Resources: League Rules and Guidelines." , Accessed: 01/02/2020.
- Major League Soccer.** 2021. "2021 MLS Roster Rules and Regulations." , Accessed: 25/01/2022.
- Mincer, J.** 1958. "Investment in Human Capital and Personal Income Distribution." *Journal of Political Economy*, 66(4): 281–302.
- Mincer, J. A.** 1974. *Schooling, Experience, and Earnings*. In , NBER Books(minc74-1): National Bureau of Economic Research, Inc.
- Oster, S. M., and D. S. Hamermesh.** 1998. "Aging and Productivity Among Economists." *The Review of Economics and Statistics*, 80(1): 154–156.
- Özdemir, A., H. Dietl, G. Rossi, and R. Simmons.** 2022. "Are workers rewarded for inconsistent performance?" *Industrial Relations: A Journal of Economy and Society*, 61(2): 137–151.
- Pallais, A.** 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review*, 104(11): 3565–3599.
- Peeters, T.** 2015. "Profit-maximizing gate revenue sharing in sports leagues." *Economic Inquiry*, 53(2): 1275—1291.
- Peeters, T., S. Szymanski, and M. Terviö.** 2022. "The Survival of Mediocre Superstars in the Labor Market." *The Journal of Law, Economics, and Organization*, 38(3): 840–888.
- Principe, F., and J. C. van Ours.** 2022. "Racial bias in newspaper ratings of professional football players." *European Economic Review*, 141.
- Prockl, F., and D. Semmelroth.** 2018. "Perception Versus Reality - Competitive Balance In Major League Soccer From 1996 – 2016." Working Papers Dissertations 36, Paderborn University, Faculty of Business Administration and Economics.
- Rosen, S.** 1981. "The Economics of Superstars." *American Economic Review*, 71(5): 845–858.
- Rupert, P., and G. Zanella.** 2015. "Revisiting wage, earnings, and hours profiles." *Journal of Monetary Economics*, 72 114–130.

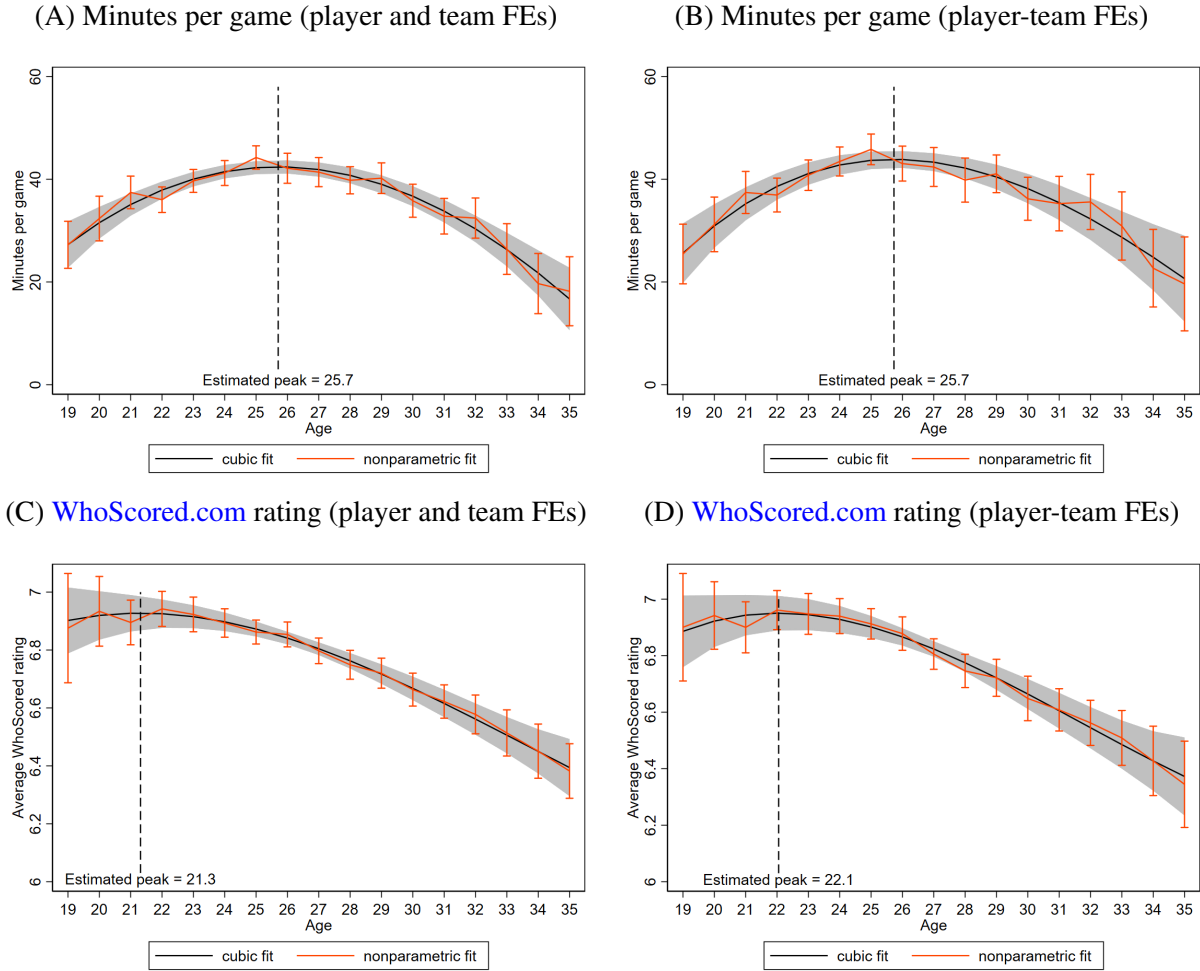
- Scarfe, R., C. Singleton, and P. Telemo.** 2020. "Do High Wage Footballers Play for High Wage Teams? The Case of Major League Soccer." *International Journal of Sport Finance*, 15(4): 177–190.
- Scarfe, R., C. Singleton, and P. Telemo.** 2021. "Extreme Wages, Performance, and Superstars in a Market for Footballers." *Industrial Relations: A Journal of Economy and Society*, 60(1): 84–118.
- Scully, G. W.** 1974. "Pay and Performance in Major League Baseball." *American Economic Review*, 64(6): 915–930.
- Simmons, R.** 2022. "Professional Labor Markets in the Journal of Sports Economics." *Journal of Sports Economics*, 23(6): 728–748.
- Solow, J. L., and A. C. Krautmann.** 2020. "Do You Get What You Pay for? Salary and Ex Ante Player Value in Major League Baseball." *Journal of Sports Economics*, 21(7): 705–722.
- Stanton, C. T., and C. Thomas.** 2016. "Landing the first job: The value of intermediaries in online hiring." *The Review of Economic Studies*, 83(2): 810–854.
- Terviö, M.** 2009. "Superstars and Mediocrities: Market Failure in the Discovery of Talent." *The Review of Economic Studies*, 76(2): 829–850.
- Twomey, J., and J. Monks.** 2011. "Monopsony and Salary Suppression: The Case of Major League Soccer in the United States." *The American Economist*, 56(1): 20–28.
- Van Ours, J. C., and L. Stoeldraijer.** 2011. "Age, Wage and Productivity in Dutch Manufacturing." *De Economist*, 159(2): 113–137.
- Walters, S. J. K., P. Allmen, and A. Krautmann.** 2017. "Risk Aversion and Wages: Evidence from the Baseball Labor Market." *Atlantic Economic Journal*, 45(3): 385–397.
- Williams, N.** 2009. "Seniority, experience, and wages in the UK." *Labour Economics*, 16(3): 272–283.
- Zangelidis, A.** 2008. "Seniority Profiles in Unionized Workplaces: Do Unions Still have the Edge?" *Oxford Bulletin of Economics and Statistics*, 70(3): 327–345.

# The Age-Wage-Productivity Puzzle: Evidence from the Careers of Top Earners

## Online Appendix

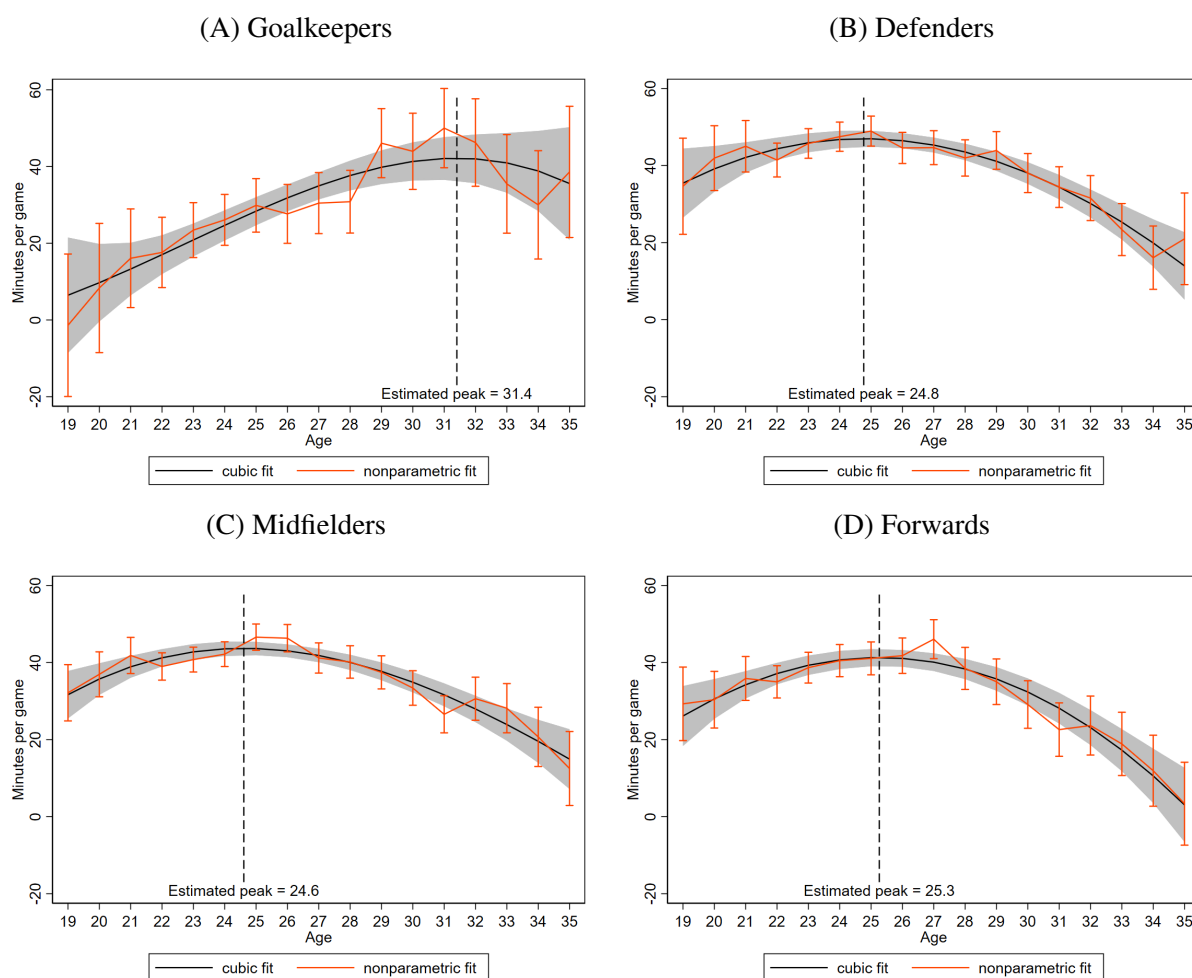
### Appendix A. Additional Figures & Tables

FIGURE A1: Estimated *age-productivity* profiles using player and team fixed effects and player-team match fixed effects specifications



Notes.- Estimates of Equations (3) & (4): minutes per game concern 2007-18; WhoScored.com ratings concern 2013-19. The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

FIGURE A2: Estimated *age-productivity* profiles using player effects: minutes per game and by playing position



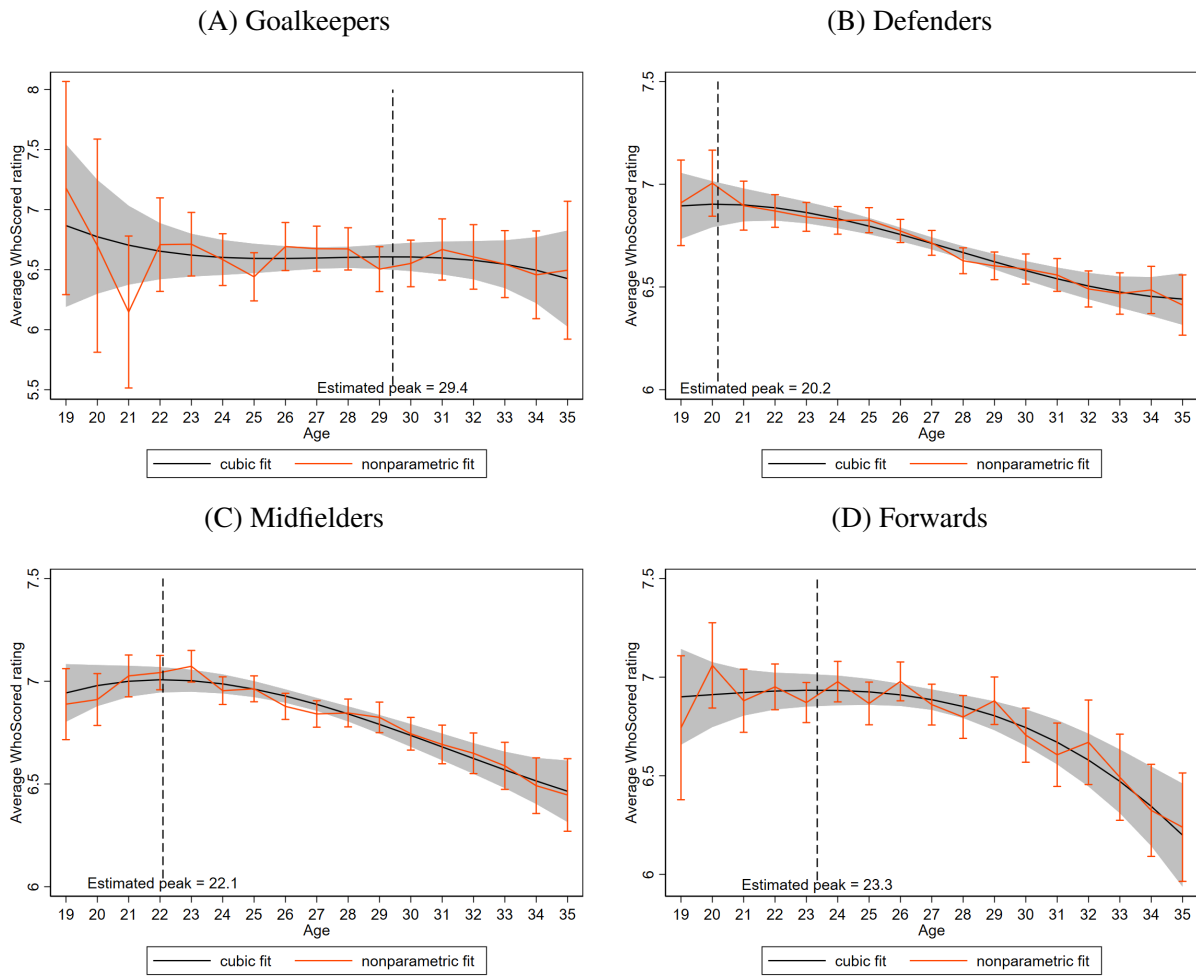
Notes.- Estimates of Equation (2) separately by position: minutes per game concern 2007-18. The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

TABLE A1: Estimated effect of team and year specific age-induced wage premium on points per game, 2007-2019

Wage premium ( $\hat{\gamma}$ )	-0.290 (0.548)
Wage bill ( $\phi$ )	2.588 (1.716)
Constant ( $\alpha$ )	-1.531 (6.577)
$N$	217
$R^2$	0.011

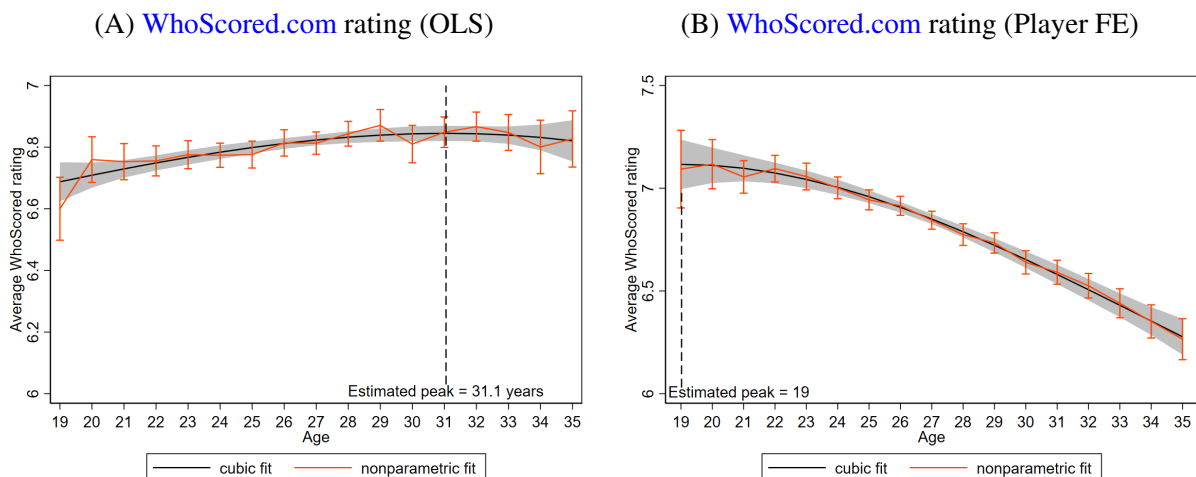
Notes.- Estimates from Equation (9). Age-wage premiums and Log(wage) are averaged at team and year level. Wages in “MLS 2019 prices” (see Section 2). Goalkeepers are excluded from the sample. Standard errors in parentheses are clustered at the team level.

FIGURE A3: Estimated *age-productivity* profiles using player effects: [WhoScored.com](#) ratings and by playing position



Notes.- Estimates of Equation (2) separately by position: [WhoScored.com](#) ratings concern 2013-19. The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

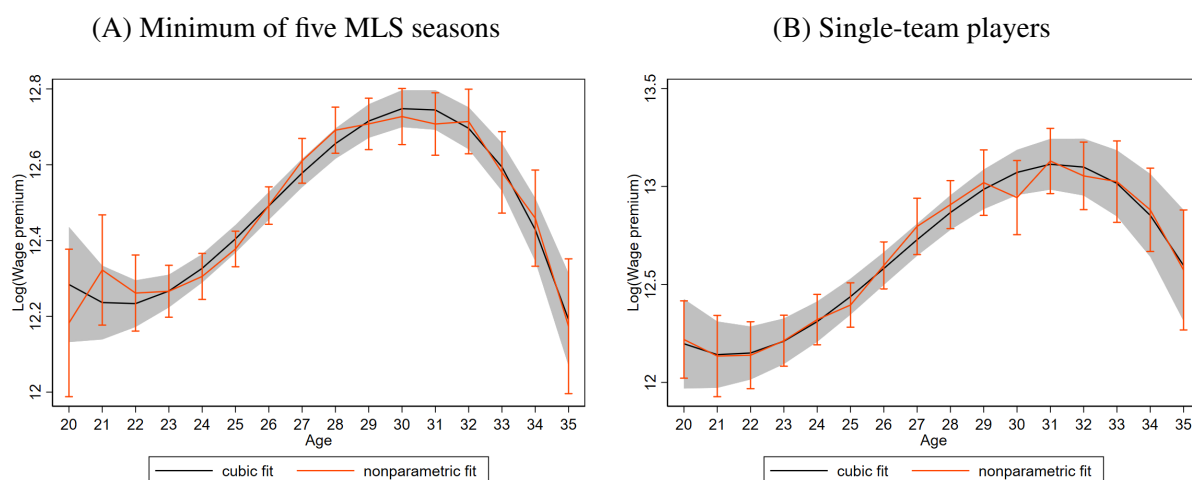
FIGURE A4: Estimated *age-productivity* profiles without de-trending



Notes.- These estimates use raw [WhoScored.com](#) rating, which have not been de-trended with respect to season and positions. The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

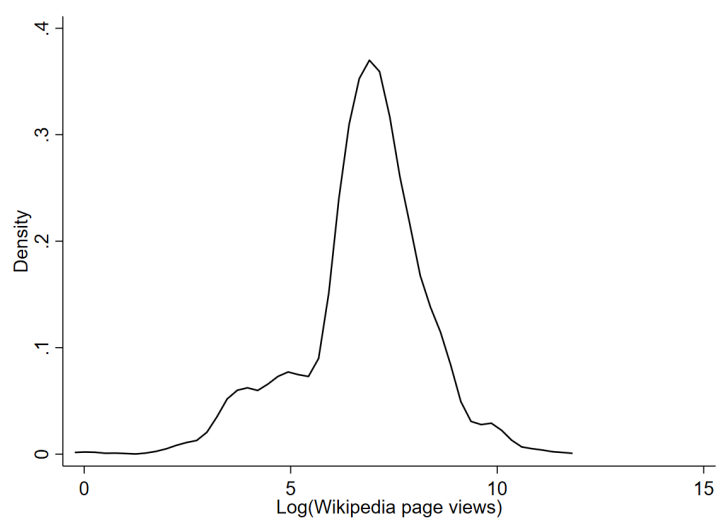


FIGURE A5: Estimated *age-wage premium* profiles using player-level fixed effects, 2008-19, outfield players only



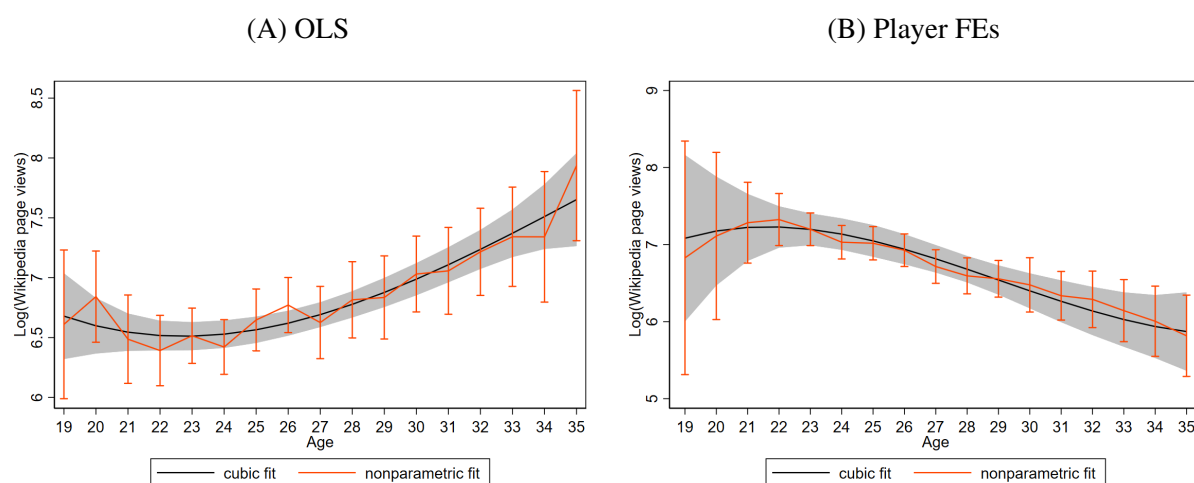
Notes.- Estimates of Equation (6): in (A) for the 367 outfield players who are observed for at least 5 seasons in MLS; in (B) for the 295 outfield players who are observed playing for only one team in MLS and at least twice with two consecutive seasons. Wages are expressed in “MLS 2019 prices” (see Section 2). The shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions.

FIGURE A6: Distribution of Log(annual Wikipedia page views) for MLS players 2016-2018



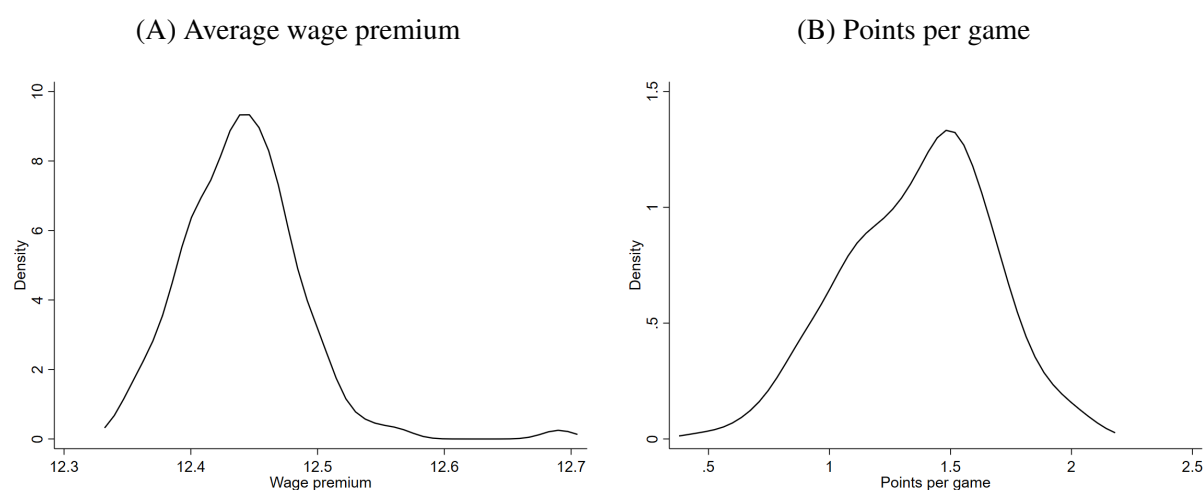
Notes.- Kernel density estimated using an Epanechnikov kernel and Silverman automatic bandwidth selection.

FIGURE A7: Estimated *age-popularity* profiles as measured by Wikipedia page views, 2016-2018



Notes.- Estimates from regressions of age on Log(Wikipedia page views). Whereas **A7A** includes no additional controls, **A7B** controls for player-specific fixed effects. Shaded area and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions. Sample contains 1,458 observations and 801 unique players.

FIGURE A8: Distribution of points per game and age-induced wage premiums across teams and years, 2007-2019



Notes.- Kernel density estimated using a Gaussian kernel and Silverman automatic bandwidth selection.

FIGURE A9: Estimated *age-productivity variance* profiles

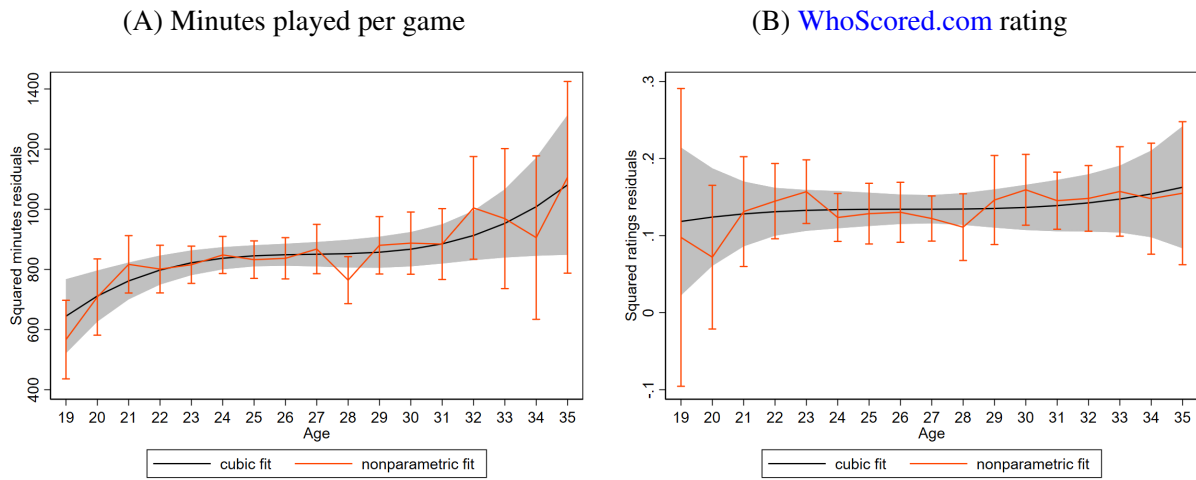
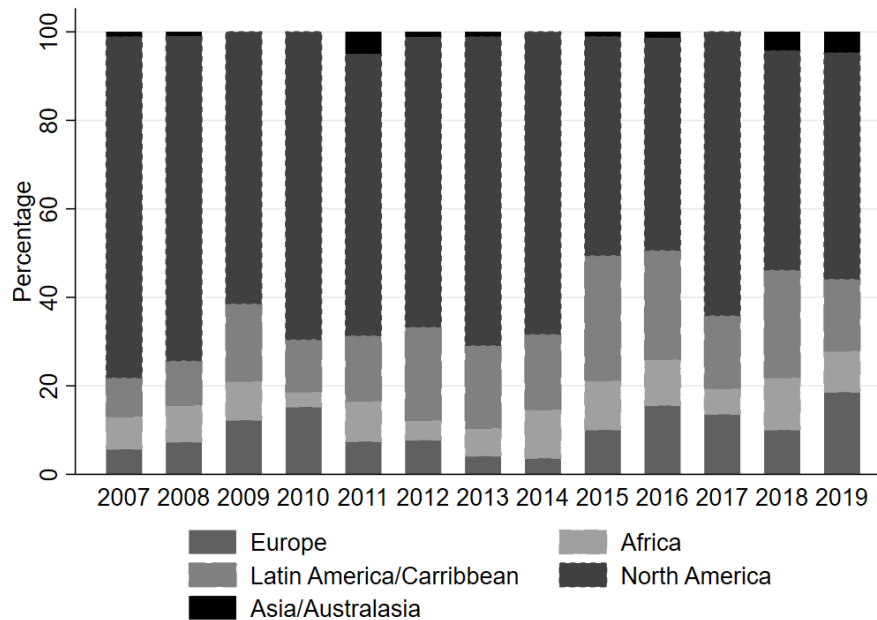


FIGURE A10: Home region of players in their first season in MLS



Notes.- Author calculations using data from MLS and Wikipedia pages. North America includes players from Mexico, which has its own professional football league that has higher television viewership in the United States than MLS (see <https://theathletic.com/4173314/2023/02/09/super-bowl-tv-ratings-soccer/>).

## Appendix B. Robustness check: Bundesliga

The main advantage of using data from MLS in this study is the availability of accurate salary data for the universe of players in the League over an extended period of time. However, there is a concern that the specificity of the league structure and wage setting in the MLS – as discussed at in Section 2 – make it hard to extrapolate the findings here onto other labour markets, or even onto other markets for professional footballers. For this reason, we conduct an additional robustness exercise where we estimate age-productivity and age-wage profiles, using the same empirical strategy as for our main analysis, but for the German top division of football; the *Bundesliga*.

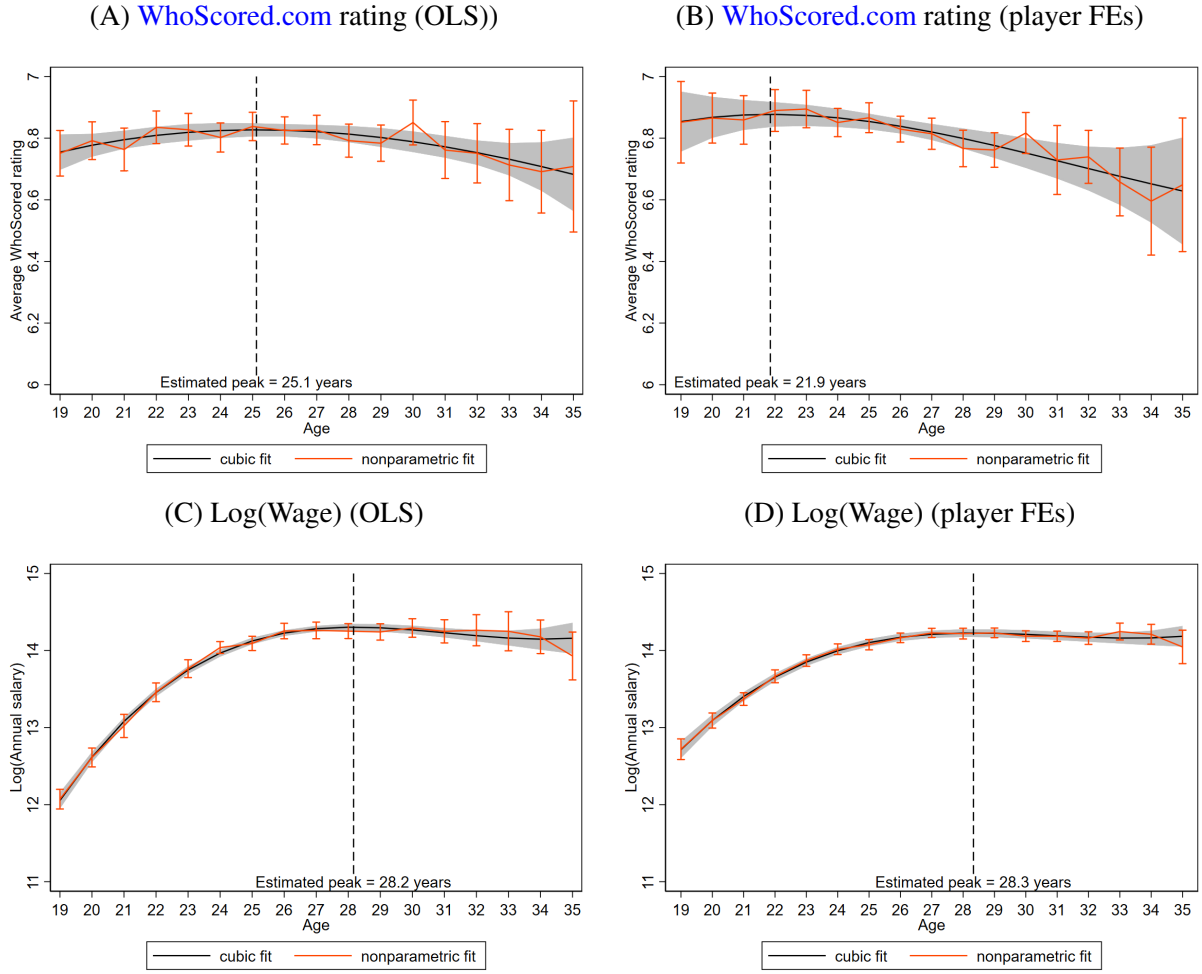
We use proxies of annual salaries or wages in the Bundesliga from the 2013/14-2022/23 seasons using data from [Capology.com](#). Since these data are not provided by the league or employers, we cannot confirm whether they are entirely accurate, and nor do they cover all the players.<sup>2</sup> However, the data are designed for use by professional clubs and agents, and have been found consistent with data from media outlets and official sources such as UEFA and FIFA reports ([Amand et al., 2023](#); [Berri et al., 2023](#)). To measure performance in Bundesliga, we use data on [WhoScored.com](#) ratings and ages for all players from the 2009/10-2019/20 seasons.

The results are largely consistent with the findings from MLS, although there are some differences. Figures B1A and B1B display the estimated age-productivity profiles. The estimated peak rating in the Bundesliga at age 21 is the same as in MLS. However, it appears that players in the Bundesliga experience less of a drop off in productivity at older ages, with 35-year-olds' ratings only estimated to have dropped 0.2 units relative to their peak, compared to a 0.5 units drop in MLS. Selection in and out of the Bundesliga seems to bias the estimates in the same direction as in MLS, with the OLS estimates suggesting an older peak age rating than the worker fixed effects model, although not to the same extent as in MLS. This is as we would expect, since the Bundesliga is a higher profile league that lacks the incentive to recruit older popular players to boost its audience. As for wages, the results in Figures B1C and B1D suggest that player wages peak at age 28, compared with age 30 in MLS. It is notable that the drop-off in wages after age 30, which is significant in the MLS, is much less apparent in the Bundesliga, where wages only drop moderately after age 30. Thus, the discrepancy between the age-productivity and age-wage profiles that we observe in MLS appears to also exist in the German top division. This adds some reassurance that our findings using the higher-quality MLS data are not completely anomalous within professional football at least.

---

<sup>2</sup>[Capology.com](#) makes the following statement regarding the accuracy of the data: “All salary figures are estimates as actual salaries may vary. Unverified players are calculated using algorithms, or sourced from reporting that has not been confirmed yet. Coverage may be limited or incomplete”.

FIGURE B1: Estimated age-wage and age-productivity profiles in the Bundesliga



Notes.- Estimates of Equation (1) and (2) using data from the German Bundesliga. Wages in “Bundesliga 2019 prices” (see Section 2). Shaded areas and orange bars represent 95% confidence intervals calculated using bootstrapping, 100 repetitions. Estimation sample sizes: 2009/10-2019/20 WhoScored.com performance (panels A and B) uses 4,029 player-season observations, with 1,304 unique players. 2013/14-2022/23 Capology.com wages (panels C and D) uses 5,248 player-season observations, with 1,876 unique players.