

Do high wage footballers play for high wage teams? The case of Major League Soccer

Article

Published Version

Scarfe, R., Singleton, C. ORCID: <https://orcid.org/0000-0001-8247-8830> and Telemo, P. (2020) Do high wage footballers play for high wage teams? The case of Major League Soccer. *International Journal of Sport Finance*, 15 (4). pp. 177-190. ISSN 1930-076X Available at <https://centaur.reading.ac.uk/92339/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Publisher: FiT Publishing

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Do High Wage Footballers Play for High Wage Teams? The Case of Major League Soccer

Rachel Scarfe¹, Carl Singleton², and Paul Telemo¹

¹ *The University of Edinburgh*

² *University of Reading*

Rachel Scarfe is currently studying for a PhD in economics at The University of Edinburgh, funded by the Economics and Social Research Centre (UK) as part of the Credit and Labour Market Foundations of the Macroeconomy Project. Her main research interests are insecure employment and labor market frictions.

Carl Singleton is lecturer (assistant professor) in economics at the University of Reading, where he is a founding member of the Football Research Group and the Sport Economics Cluster as well as cofounder of the Reading Online Sport Economics Seminar Series. He completed his PhD at The University of Edinburgh in 2017.

Paul Telemo is currently studying for a PhD in economics at The University of Edinburgh, funded by the Economics and Social Research Centre (UK). His main research interests are social mobility and labor economics.

We are grateful for helpful comments from three anonymous referees, Dennis Coates, J. James Reade, and Daniel Schaefer, as well as for research assistance from Leon Singleton. The data used in this study are publicly available from Major League Soccer (MLS) LLC and the MLS Player's Association.

Singleton thanks the Economic and Social Research Council (UK) for funding support under Grant No. ES/J500136/1.

Abstract

Intuition and sports knowledge suggest that the most talented professional footballers play for the best teams, i.e., positive assortative matching based on productivity. We consider Major League Soccer (MLS) between 2007 and 2017. We estimate a wage equation, finding that player and team fixed wage premiums are negatively correlated. This is a puzzle, especially because our estimates of players' wage premiums do correlate strongly with measures of their performance on the pitch, and there is evidence of positive teammate sorting. However, the estimated wage premiums of MLS teams are highly and negatively correlated with their success in the league and their home game attendance. This finding is consistent with an explanation whereby part of an MLS team's success comes from its ability to bargain down the price that it pays for talent.

Keywords: firm-specific wages, AKM wage equation, matching, superstar pay

JEL codes: J31, J49, Z22

<http://doi.org/10.32731/IJSF.154.112020.03>

Introduction

The question of whether high productivity workers are matched to high productivity firms is of considerable interest, not least because it sheds light on the matching process taking place between workers and firms (e.g., Becker, 1973; Shimer & Smith, 2000). To address this, studies have used longitudinal linked employer-employee data to estimate AKM (Abowd, Kramarz, & Margolis, 1999) wage equations, which admit both worker and firm fixed effects: If the fixed effects are indicative of worker and firm productivity, then the sign and magnitude of the correlation between these two sets of estimated effects may indicate whether a labor market features assortative matching of worker and firm productivity types (e.g., Andrews et al., 2012; see Lopes de Melo, 2018, for a summary of the evidence on assortative matching in labor markets using the AKM methodology).

Association football, or soccer, has been extensively studied by labor and sports economists (Frick, 2007). Due to the global popularity of football and the nature of the game, there is extensive data on measures of worker (player) and firm (team) productivity. Data on player wages are rarer, but where they exist, they have been combined with readily available productivity data to attempt to draw insights about labor markets in general. Because the market for football talent

features high levels and high variance of earnings, much of this literature has focused on the impacts of superstars and what determines their wages (e.g., Lucifora & Simmons, 2003; Lawson et al., 2008; Coates et al., 2016; Jewell, 2017) as well as the related migrant wage premium (e.g., Bryson et al., 2014). Other studies have more generally asked what determines the pay of football players (e.g., Kuethe & Motamed, 2010; Thrane, 2019).

The fact that the market for professional football players is generally closed off from the rest of the labor market, with very little mobility in or out during the typical career span of a footballer, makes this an appropriate context to study the assortative matching of workers and firms. To investigate this, we estimate an AKM wage equation using data on the earnings of professional football players in US Major League Soccer (MLS) between 2007 and 2017. Despite the very particular nature of this labor market—including its rules and regulations, which we elaborate on below—there are some advantages of studying it, reflected in the number of previous studies that have done so (e.g., Kuethe & Motamed, 2010; Coates et al., 2016; Jewell, 2017). First, we observe the salaries of almost all players rather than a limited sample. Second, our data include worker and firm productivity measures (e.g., minutes played and League performance), which means that we can sense check the interpretation of the results from applying the AKM methodology to this question. We would be naïve to think that our estimates from an AKM wage equation in this context provide conclusive evidence on the presence (or not) of assortative matching on productivity. Nonetheless, we can answer the more straightforward question of whether high wage footballers play for high wage teams, as fundamentally, despite the peculiarities, this is a market where players and teams negotiate pay, where players receive and reject offers from teams, and where teams search for and match with talent.

We find that the correlation between estimated player and team fixed wage effects is small and negative. This result is robust to varying the sample selection, such as excluding the superstar designated players in MLS (e.g., David Beckham). High wage footballers do not appear to play for high wage teams in this talent market. Yet intuitively, and based on sports knowledge, we would expect that the most talented footballers do in fact play for the most successful teams. As suggestive evidence that this might nonetheless be the case in MLS, there is evidence of player segregation across teams, with players earning high wage premiums tending to have teammates with similarly high wage premiums.

MLS is unique when compared with other professional football leagues. The League itself is closed, with no promotion or relegation, and the only change in membership is through franchise expansion or dissolution. There are detailed regulations concerning pay and the mobility of players in MLS (see Kuethe & Motamed, 2010, or Coates et al., 2016, for a comprehensive discussion). Like other US sports leagues, there is a salary cap, but it only covers the first 18–20 of up to 30 players registered on a team's roster for the season.¹ Most players in MLS are technically not paid by their team but by the League itself, with which they have a contract. Despite this, teams negotiate with players on their rosters to determine salaries. The teams' overall budgets affecting these wage negotiations, besides what appears in some sense to be a nominal salary cap, are determined by their revenues, mostly from attendance at home games (Brownlee & Lorgnier, 2017). Teams can also hire so-called designated players, between one and three each season over the period we study, whose wages are mostly outside the salary cap and are paid directly by the teams themselves rather than the League. In most cases, these are the highest paid players in a team. Each designated player contributes to the salary cap up to a fixed percentage of the overall cap level. Teams pay a lump sum tax to the League for each designated player. To level the playing field, the League also makes not insignificant monetary allocations to some teams each season, typically those that do not make the end of season playoffs through their performances on the pitch, effectively raising their salary caps (Coates et al., 2016). Similarly, if the teams fill their third designated player slot, which the higher revenue teams are most likely to do, the tax to the League from doing so is redistributed to the others above the cap, raising their purchasing power in the market.

Due to the peculiar contracting situation in this talent market, player mobility between teams is to an extent controlled by the League under the roster regulations mentioned above. That said, players and teams are not forced upon one another. Players can reject moves and offers made by other teams. They also have outside options, to varying degrees, both inside and outside of MLS. Similarly, teams search for, evaluate, and choose the players that make up their rosters. Combined with the fact that players and teams negotiate pay, the question of what the general sorting and matching patterns are between employer and employee in this labor market are just as relevant as in any other. However, the peculiarities of MLS make it not at all obvious that these patterns will be similar to those found elsewhere.

This study is not the first to consider the issue of assortative matching in sport, or even in football. Gandelman (2008) showed that positive assortative matching exists in football, in terms of players' performances on the pitch being correlated with the nature of their mobility across teams, with good performances leading to promotion to a better team and poor performance leading to demotion. Outside of football, Filippin and van Ours (2015) found similar patterns by studying the performance and mobility of elite long-distance runners across teams in Italy. Returning to football, Drut and Duhautois (2017) used player and team performance data from five seasons of Italian professional football, along with wages, to find substantial evidence of positive assortative matching using an AKM approach. They estimated a correlation between player and team fixed wage premiums of 0.46. Because their data featured a high degree of mobility and the approximate universe of players over their sample period, they argued that the AKM-type wage models are appropriate for detecting assortative matching when the amount of mobility observed is not limited. In our study, we apply a similar approach to Drut and Duhautois (2017), but we study a longer sample period of 11 seasons. This ought to allow more precise and robust estimates of the AKM wage equation as well as feature greater mobility of players between clubs with which to identify relative team wage premiums. Furthermore, the MLS is a very different institution when compared with Italian professional football, not least in terms of competitive balance: In Italy, the same team, Juventus FC, have been the champions in each of the past eight seasons between 2011/2012 and 2018/2019, whereas over the same period only one team won the MLS Cup more than once.

Like Drut and Duhautois (2017) showed for Italy, we find that the players' estimated wage premiums in MLS do correlate strongly with measures of their productivity for teams on the football pitch. However, unlike in Drut and Duhautois (2017), this is not the case for the estimated team wage premiums: More successful teams, both in terms of winning football matches and attracting supporters, and thus revenue, pay significantly lower wages in MLS.

There are potential ways to reconcile the particular differences between our results for MLS compared with the findings from Drut and Duhautois (2017) for Italy. The first is negative limited mobility bias in the estimated correlation of the player and team fixed effects (Andrews et al., 2008). However, like Drut and Duhautois (2017), the data we use are an approximate universe of players and teams, and we show this source of bias is not an issue. Second, it may be difficult to identify assortative matching using wage data alone. A small or negative correlation estimate from an AKM wage equation does not rule out that, in fact, high productivity workers move to and match with high productivity firms (e.g., Eeckhout & Kircher, 2011; Lopes de Melo, 2018; Bartolucci et al., 2018). There are specific reasons why this might be the case in a football talent market, all of which plausibly result in higher productivity (more successful or richer) teams having relatively greater bargaining power with players and thus being able to extract a greater share of the surplus from successfully matching with a player. First, salary regulation could affect hiring decisions and distort the distribution of surplus when a player and team match. Second, there are well-known examples throughout football of players who turn down lucrative wage offers from more successful teams to instead play for teams for which they have some personal affinity, such as their home-town team.² Closely related, footballers perhaps have some preference to play for productive (successful) and well-supported teams, and their decisions might not be just about the money offered directly by teams. In other words, players who switch to more successful teams may be willing to accept relatively low wages. Third, some teams may be successful in a league simply because they are better able than others to apply their monopsony power over players to bargain down wages. This latter potential explanation is particularly relevant here, given the peculiar institutions governing the MLS talent market, not least the salary cap.

The remainder of the paper proceeds as follows: the next section introduces our data, the one following describes the Mincerian AKM wage equation that we estimate, then we present and discuss our results, followed by our conclusions.

Data

The data on player wages and team affiliations come from the MLS Players Association (MLSPA), the collective bargaining representative of players during the negotiation of league-wide salary rules. These data refer to midway through each MLS season in August of each year, after the secondary transfer window when players can be signed from abroad. The wage measure is the guaranteed annualized compensation or salary, henceforth referred to as wages, which includes payments for signing with a team or related to marketing.³




Figure 1. Percentiles of MLS Player Log (Guaranteed) Wages, 2007–2017

Notes: Author calculations using MLSPA data. From bottom to top, shows 10th, 25th, 50th, 75th, and 90th percentiles.




Figure 2. Team Guaranteed Salary Totals for the 2017 MLS Season

Source: Major League Soccer Players Association.

Figure 1 tracks percentiles of player log wages between 2007 and 2017. This was a period of rapid pay growth in MLS, with nominal wages more than doubling across the distribution and an increase in the median of 160%. This increase coincided with the growing popularity of the league, higher revenues, and more foreign players choosing to play in the US (Brownlee & Lorgnier, 2017). The league expanded in this period from 13 teams in 2007 to 22 in 2017.⁴

Despite the salary cap and other regulations mentioned earlier, it appears that teams in MLS are operating on a salary budget rather than an effective league-wide cap. For example, Figure 2 shows the guaranteed salary totals for teams in the 2017 season. There is substantial variation across teams, ranging from \$5 million for the lowest budget teams to over \$20 million for the highest, whereas the salary cap for roster spots 1–20 that year was just \$3.85 million.

We obtain information on a player’s age, the number of games they started, minutes played during a season, designated player status, assists, goals and shots at goal, nationality, and team-level performance data, including season average attendances, from the official MLS website.⁵ While the time-varying individual performance statistics (goals assists, etc.) are not used in our wage regression models, we will later use them to verify whether our estimates of player fixed wage effects reflect some measures of player productivity. A handful of player observations are dropped from the dataset where we could not reliably match a record in the salary data with the information on player characteristics. In our preferred specification and results, we exclude any observations where the player did not play for their team in that season, thus ignoring any players who were either continuously injured or just making up the numbers on a team’s roster. This leaves 4,213 player-year observations, covering 1,410 different players. The minimum, median, and maximum ages of players in the analysis are 16, 25, and 42, respectively. Almost 7% of player-year observations are designated players. The median number of teams played for over player-year observations is 2, implying a significant amount of mobility between teams. Just under 30% of the players represented in the main analysis played for more than one MLS team in 2007–2017. Although the movement of players between teams is not as free as in European football leagues, or in labor markets more generally, switches between teams do happen frequently. There are currently twelve different ways listed by MLS through which a team can acquire a player, most of which can then be observed as the intra-league movements of players.⁶ Further descriptive statistics for the main analysis sample of players can be found in Table 1.

Method

To decompose wages into components associated with players and with teams, we use an AKM wage equation with player and team fixed effects. Let there be $i = 1, \dots, P$ players, $j = 1, \dots, J$ teams, $t = 1, \dots, T$ years, T_i years per player, and $N = \sum_i T_i$ total observations. The standard Mincerian AKM-type wage equation is given by

$$w_{it} = \alpha_i + \varphi_{j(it)} + \mathbf{x}_{it}^T \boldsymbol{\beta} + \delta_t + \varepsilon_{it} \tag{1}$$

Table 1. Descriptive Statistics for MLS Player Analysis Sample, 2007–2017

Player-level descriptive statistics					
	Mean	SD	Min.	Max.	N
Seasons played	2.99	2.41	1	11	1410
Team changes	0.46	0.89	0	6	1410
From USA/Canada (%)	0.47				1410
Goalkeeper (%)	0.08				1410
Defender (%)	0.34				1410
Midfielder (%)	0.34				1410
Forward (%)	0.24				1410
Full-sample (player × season) descriptive statistics					
	Mean	SD	Min.	Max.	N
Age	26.57	4.32	16	42	4213
Salary (\$)	243,894	648,969	12,900	8,650,000	4213
Designated player (%)	0.07				4213
Goals per 90 min	0.12	0.18	0.00	1.64	3837
Assists per 90 min	0.12	0.14	0.00	1.34	3837
Shots per 90 min	1.25	1.05	0.00	6.59	3837
Games played	19.27	9.32	1	35	4213
Minutes played	1398.47	852.56	14	3060	4213

Notes: Author calculations using MLS and MLSPA data. Goals and assists are not included for goalkeepers

where w_{it} is the log wage for player i in year t and α_i is the player fixed effect. Team fixed effects with mean zero are given by $\varphi_{j(it)}$, where $J(it) = j$ indicates the team employing player i in year t . In this specification, the player and team fixed effects capture constant unobserved heterogeneity in players and teams, respectively. \mathbf{x}_{it} is a vector of time-varying observable characteristics: a dummy variable indicating designated players—player age minus 15, to roughly represent football career experience—and this value squared and cubed, all to address football-life-cycle wage growth. δ_y gives year effects, accounting for average annual wage growth in MLS. The remaining heterogeneity in wages is in the residual, ε_{it} .

Equation (1) can be estimated using least squares with a strict exogeneity assumption: $E[\varepsilon_{it} | \mathbf{x}_{it}, \alpha_i, \varphi_{j(it)}, \delta_t] = 0$.

In words, the estimated team fixed effects will be biased if players change teams due to some component of ε_{it} , such as transitory shocks to team-wide or player-team match-specific wage effects. This is equivalent to assuming that players are assigned randomly to teams, conditional on the observable characteristics \mathbf{x}_{it} and δ_t and the unobservable time-invariant heterogeneity captured by α_i and $\varphi_{j(it)}$. Most applications have found that labor market data appear to support this exogenous mobility assumption (see Card et al., 2018, for a general discussion of this assumption in reality and practice). As a rough test for one implication of this assumption, we look at the year-to-year changes in residual wages for players at their original teams prior to moves within MLS. We obtain measures of residual wages by estimating the equivalent of Equation 1 without the player or team fixed effects; i.e., we estimate using least squares the equation

$$w_{it} = \mathbf{x}_{it}^T \tilde{\boldsymbol{\beta}} + \delta_t + \eta_{it} \tag{2}$$

where $\eta_{it} = \alpha_i + \varphi_{j(it)} + \varepsilon_{it}$ is a residual measure of a player’s wages after controlling for their experience in MLS and average annual wage growth. If players move between teams exogenously, then on average we should not observe any growth or decline in their residual wages prior to a player move. For the 321 player moves in our main analysis sample where that player had played for their previous team for at least two years, we find no significant evidence that mobility is related to a change in wages at the previous team. On average, these players experienced annual residual wage growth of 2.6% between the two seasons before their moves, with a median value of 0.8%. The mean residual wage growth is significantly different from zero and non-negative, providing some evidence that mobility is not related to depressed wage growth at the previous team.

The player fixed effects are constant and transferable, affecting wages wherever and whenever a player features in MLS throughout their career in the league. The team fixed effects are wage premiums received by all players within a team, identified by the players who switch teams. Therefore, estimates of these premiums are only comparable within connected sets of players and teams (Abowd et al., 2002). As we observe the salaries of almost all players during their time in MLS, in our main results, all of the players and the 23 teams are connected in a single set, meaning our estimates will be fully representative of this particular labor market in this period. It is worth noting that we are only estimating the wage premiums for these individuals in their particular occupation as footballers in MLS, which is just a portion of their working lives. For instance, the model is completely silent on what wage premiums these particular workers gain from playing football in MLS as opposed to some other occupation or even how the wage premiums in MLS compare to other football leagues where these players may find themselves during their careers. For robustness, we later vary the dependent variable of the wage regression, considering the log of wages per minute played. We also estimate Equation 1 excluding some players, such as designated players, who do not contribute fully to the teams’ salary caps, and non-domestic players.

To assess whether high wage players generally play for high wage teams in MLS, we consider the correlation between the estimated player and team fixed effects from Equation 1, $Corr_{it}(\hat{\alpha}_i, \hat{\varphi}_{j(it)})$. This measure is negatively biased because the fixed effects are estimated with errors (Andrews et al., 2008). This limited mobility bias is large in small sample settings and related to the amount of observed mobility between teams. Andrews et al. (2008, 2012) predicted and demonstrated empirically that the estimated value of $Corr_{it}(\hat{\alpha}_i, \hat{\varphi}_{j(it)})$ will be increasing and concave in the number of observed movers per team, asymptoting toward its true value as this mobility increases.

Results

Column 1 of Table 2 shows our preferred estimates of Equation 1 (Columns 2–5 present the results of robustness checks, which are discussed later). The fit of the wage equation is high, with $R^2 = 0.91$. This fit increases only slightly to 0.95 when we estimate a version with player-team (match-specific) fixed effects rather than separate player and team effects,

Table 2. AKM-Style Regression Model Estimation Results for Log Player Wages, 2007–2017

	(1) (Main)	(2) (All players)	(3) (No DPs)	(4) (per 90 mins)	(5) (> 2y)
Designated player	0.652*** (0.109)	0.658*** (0.108)		0.536*** (0.108)	0.655*** (0.112)
Age (years)	-0.070 (0.053)	-0.124*** (0.041)	-0.073 (0.052)	-0.411*** (0.084)	0.093 (0.059)
Age squared	0.012** (0.005)	0.016*** (0.004)	0.012** (0.005)	0.030*** (0.007)	0.013** (0.005)
Age cubed	-0.0006*** (0.0001)	-0.0006*** (0.0001)	-0.0005*** (0.0001)	-0.0008*** (0.0001)	-0.0005*** (0.0001)
Player, team, & year effects	Yes	Yes	Yes	Yes	Yes
R ²	0.913	0.922	0.862	0.685	0.849
RMSE	0.349	0.342	0.341	0.732	0.361
St. dev. of log wages: $Std(w_{it})$	0.959	0.981	0.743	1.057	0.929
St. dev. player effects: $Std(\hat{\alpha}_i)$	0.893	0.909	0.765	0.965	0.723
St. dev. team effects: $Std(\hat{\phi}_{j(it)})$	0.092	0.091	0.089	0.146	0.097
St. dev. observables: $Std(\mathbf{x}_{it}^T \hat{\beta})$	0.601	0.591	0.551	0.728	0.673
Correlation: $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$	-0.102	-0.097	-0.084	-0.121	-0.110
N: player-years	4,213	5,195	3,906	4,213	3,138
P: players	1,410	1,820	1,320	1,140	606
J: teams	23	23	23	23	23
M/J: mobility	27	30	26	27	25

Notes: ***, **, and * indicate significance from one at 1%, 5%, and 10% levels, respectively, two-sided tests. Values in parentheses give standard errors robust to player-level clusters. Dependent variable is the log of guaranteed salary (\$), except for column (4), where it is the log of guaranteed salary per 90 minutes played during the season.

suggesting that the assumed additive separability in Equation 1 is a workable assumption. In other words, the extent to which there are synergies between players and teams that are reflected in players’ wages is small. Unsurprisingly, the designated player premium is large and significant at 65 log points, which is identified from the players who switch between being a designated player or a standard roster player. The estimated coefficients on the age terms imply that, on average, a player’s wages fall up to age 20, then increase and peak at a 2.5% premium by age 27, thereafter declining rapidly, with a player age 35 earning just 55% of what he would have earned at 29.⁷

The estimated team fixed effects, $\hat{\phi}_j$, are shown in Figure 3. Relative to the average player-year wage, these team wage premiums range from -23 log points to 18 log points. Despite what might appear to be sizeable premiums, these are dwarfed by the overall wage variation in MLS.⁸ Individually, these wage premiums are not significantly different from zero at the 95% level. However, pair-wise, there are significant differences, implying that this talent market does feature team-specific wage premiums to some extent. The naïve AKM-implied measure of assortative matching, the estimated correlation between the player and team fixed effects, $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$, is -0.10. In this football talent market, high wage footballers (α_i) do not play for high wage teams (ϕ_j).

To interpret this result, we first look more closely at our estimated player fixed effects. If the estimates of α_i capture something of the productivity of each player, then we would expect them to correlate positively with player-level measures of productivity. In Figure 4(A–D), we look at the relationship between four measures of individual player productivity and the estimates of α_i . Panel A shows that a player’s share of his team’s total minutes played during a season is positively correlated with the estimated wage effects, and we might assume that more productive players on a team’s roster play more often during a season. Panel B looks at players whose primary position was in midfield and for whom a suitable measure




Figure 3. Estimated Relative Team Log Wage Premiums, $\hat{\varphi}_j$

Notes: Estimates relate to column 1 of Table 2. By construction, the average wage premium over player-year observations is zero. The dashed lines show 95% confidence intervals robust to player clusters.

of productivity would reflect their creativity, such as the number of goal assists they generate per 90 minutes. This is positively correlated with the player wage effects, with a correlation coefficient of 0.26, which would potentially be higher if we were able to filter out more defensive-minded midfield players. The primary aim of a forward in football is to score goals. Panels C and D show that the number of goals that forwards score and the number of shots at goal that forwards take, both per 90 minutes played, are strongly positively correlated with their estimated fixed wage effects.

The evidence above suggests that the estimates of α_i reflect individual player productivity. Similarly, the team fixed effects estimated from the AKM-type wage model could be expected to reflect some type of overall team-level productivity that is not specific to any one player. In MLS, this could derive from the team’s support structure for its players, i.e., the facilities or the management. We consider two possible observable measures of a team’s productivity: the number of regular season points achieved per game and the average home attendance.⁹ There is substantial literature that relates the two measures in sport, as well as in MLS specifically (e.g., Jewell, 2017; Wooten, 2018), but home attendance should also capture the talent purchasing power of teams, as this is their major source of revenue.¹⁰ We are able to compare the estimated player fixed effects with these team-level productivity measures. Given the negative correlation we observe between $Corr_{it}(\hat{\alpha}_i, \hat{\varphi}_{J(it)})$, we would not necessarily expect our estimates of α_i to correlate positively with overall measures of productivity. Panel E of Figure 4 plots each player’s estimated fixed wage effect against the average over seasons of the points per game achieved by the team they play for each season. The correlation between the two statistics is slightly positive but weak at just 0.084. Similarly, in Panel F, the correlation between each player’s estimated fixed wage effect and their teams’ average home attendance is practically zero.

Our prior assumption, based both on economic theory and sports knowledge, would be that the most productive soccer teams pay the highest wages, which ought to be reflected in the estimates of the team fixed effect, φ_j . Figure 5(A) shows that there is a negative correlation between the average over seasons of the league points achieved per game and the estimated team fixed effects. The Pearson correlation is -0.52, suggesting teams that pay a high wage premium to all players tend to have less success in the league. Similarly, Panel B shows that the negative correlation between average home attendance and the estimates of φ_j is even stronger than for points per game, which rather counter-intuitively suggests that the better supported teams, with larger stadiums and revenues, pay lower wage premiums than the less well-supported teams. The allocation monies distributed by MLS to less successful teams may partially account for this result, though they do not generally compensate entirely for differences between teams in their ability to generate revenue.

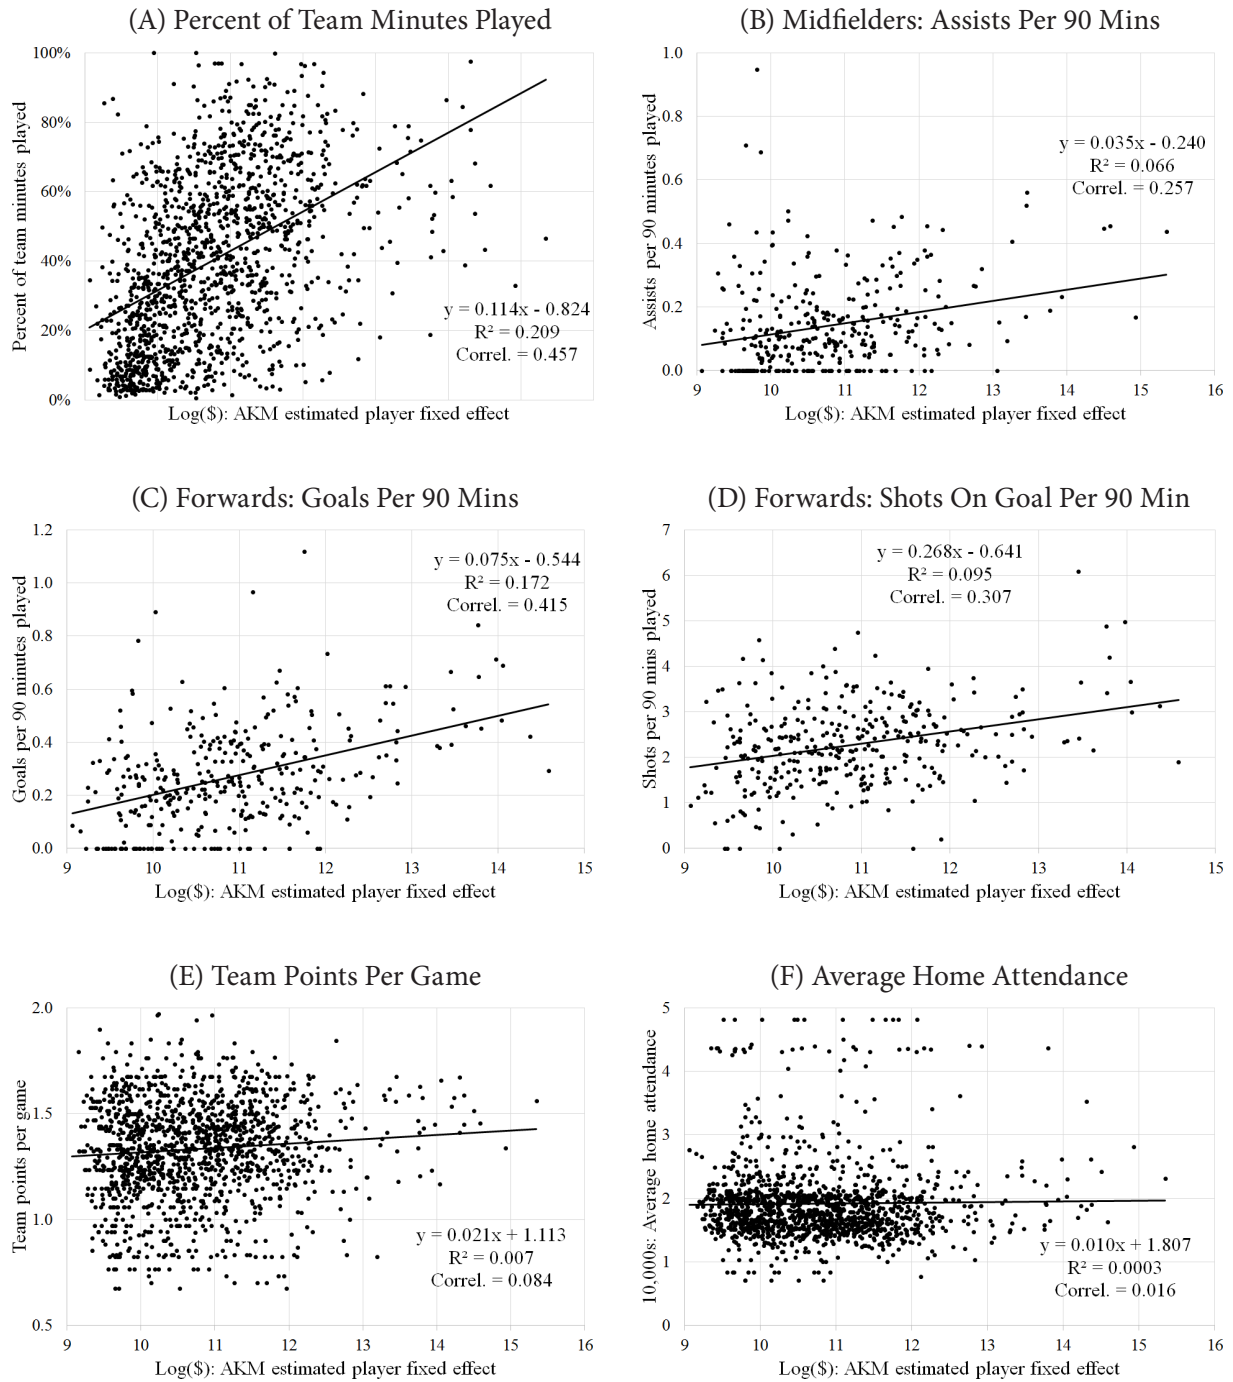


Figure 4. Correlation of Estimated Player Fixed Wage Effects with Selected Player Performance Statistics, Their Teams' Performance, and Average Home Attendance

Notes: Estimates relate to column 1 of Table 2. For each player, performance and team-level statistics are first computed within season, and then an average is taken over the years each player appears in MLS. In panels A, E, and F, there are 1,410 players represented. There are 573 players in panel B whose primary position was midfield, and 343 players in panels C and D whose primary position was forward.

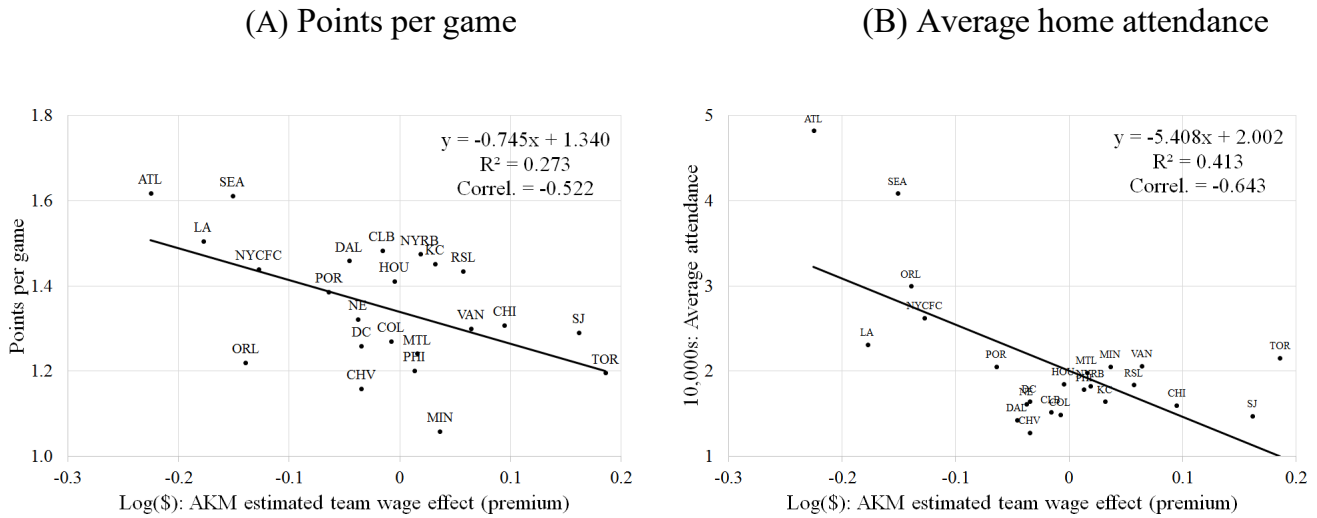


Figure 5. Correlation of Estimated Team Fixed Wage Effects with their Performance and Average Attendance
 Notes: Estimates relate to column 1 of Table 2. Averages were taken of each year’s value over seasons when the team appeared in MLS between 2007–2017.

Table 3. AKM-Style Regression Model Estimation Results for Log Player Wages, 2007–2017: Additional Robustness Checks

	(1) (> 8y)	(2) (USA/CAN)	(3) (Foreign)	(4) (Year > 2010)	(5) (Year ≤ 2010)
Designated player	0.896*** (0.084)	0.979*** (0.084)	0.436*** (0.062)	0.543 (0.401)	0.301*** (0.058)
Age (years)	0.126 (0.083)	-0.109 (0.060)	0.273 (0.478)	0.230 (0.194)	0.036 (0.056)
Age squared	0.008 (0.007)	0.015*** (0.003)	0.004 (0.005)	0.010 (0.015)	0.010* (0.005)
Age cubed	-0.000* (0.000)	-0.001*** (0.000)	-0.000** (0.000)	-0.001 (0.000)	-0.000*** (0.000)
Player, team, & year effects	Yes	Yes	Yes	Yes	Yes
R ²	0.815	0.871	0.946	0.937	0.940
RMSE	0.364	0.354	0.325	0.325	0.278
St. dev. of log wages: $Std(w_{it})$	0.609	0.691	1.107	0.731	0.780
St. dev. player effects: $Std(\hat{\alpha}_i)$	0.333	0.612	0.943	0.539	0.230
St. dev. team effects: $Std(\hat{\phi}_J(it))$	0.018	0.010	0.040	0.040	0.006
St. dev. observables: $Std(\mathbf{x}_{it}^T \hat{\beta})$	0.355	0.334	0.798	0.621	0.318
Correlation: $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_J(it))$	-0.222	-0.068	-0.216	-0.208	-0.173
<i>N</i> : player-years	859	2,432	1,781	655	1,853
<i>P</i> : players	93	666	744	569	1,136
<i>J</i> : teams	23	23	23	16	23
<i>M/J</i> : mobility	7	18	9	9	17

Notes: ***, **, and * indicate significance from one at 1%, 5%, and 10% levels, respectively, two-sided tests. Values in parentheses give standard errors robust to player-level clusters. Dependent variable is the log of guaranteed salary (\$).

Robustness Checks

We perform two sets of robustness checks. Firstly, despite using the approximate universe of players in this period, we nonetheless check whether limited mobility in the sample is significantly biasing the estimated value of $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$ downward, as shown to occur when estimating this type of wage equation by Andrews et al. (2008). If this were the case, then it might indicate that there is in fact a positive correlation between player and team fixed effects and that high wage footballers do indeed play for high wage teams. To check this, we drop at random 10%, 20%, . . . 90% of the observations and reestimate Equation 1 for the remaining largest connected sets of players and teams. In Figure 6, we plot the estimated correlation between the player and team fixed wage effects from each of these smaller samples against the number of observed player moves per team, M/J . As per Andrews et al. (2012), in smaller samples of MLS players, we find that the negative bias is substantial and decreasing with the amount of mobility, M/J , asymptoting to the figure of approximately -0.1 when we increase the sample size.

Secondly, we check that our results are not driven by a small group of particular players or by a particular period in MLS. To do so, we estimate several alternative versions of Equation 1, each time changing only one aspect of the wage model or player sample compared with the main results. In Column 2 of Table 2, we show results using all player-year observations rather than only those who played for their teams during a season. This increases the number of player-year observations by almost 20% and adds 410 players, or 30%, reflecting that these players do not tend to stay in MLS as long as those who are playing in games. There is very little change in our estimate of $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$, suggesting that our results hold regardless of how much time players spend on the pitch.

Column 3 of Table 2 drops designated player-year observations from the sample used in the preferred results, including the 90 players who held designated status for their whole MLS careers. Although we address an average designated player wage premium in our preferred results of column 1, we nonetheless check whether our results are affected by any high wage outliers from the mega superstars of MLS (e.g., David Beckham, Kaká, Steven Gerrard, David Villa, or Thierry Henry). There are three reasons why our results might be sensitive to the inclusion of designated players in our sample. First, designated players are not subject to the same salary regulations as other players in the league. Second, they are less mobile within MLS than other players. Third, it is likely that a greater proportion of their salary comes from product endorsements and corporate sponsorship.¹¹ If this is true, then these players may be less concerned about earning a high wage when choosing to play for a high productivity team. As we would expect, the standard deviation of both log wages and player fixed effects is lower in this specification. However, our estimate of $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$ is similar to that from our original specification. This suggests that the negative correlation we find is not due to a small number of designated players.

Column 4 of Table 2 uses the preferred sample of column 1 but instead has a dependent variable of the log of wages per minute played during the season. Because the wages in our data are annual salaries, this specification estimates player and team fixed wage effects, which are more likely to be strongly correlated with innate values of realized productivity on the football pitch. Column 5 of Table 2 drops from the analysis any players who featured in MLS for less than three years of the sample period to focus on a subset of players who will generally have higher mobility. This drops 804 player observations, or 57%, but just 26% of the player-year observations. Across all these different specifications, the main result on $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$ remains robust, with the estimate varying between -0.08 and -0.12. Therefore, the estimates are not sensitive to the superstars in MLS, nor are they driven by the mobility of players who actually play very little or by players who leave MLS after only a short period.

Table 3 presents the results of further robustness checks, which represent more substantial changes to our estimation approach when compared with the main set of results. Column 1 only considers players with more than eight years of experience in MLS during the sample period. These players would be eligible for free agency following its introduction from 2016 onwards, i.e., they then had the right to negotiate with any team in the league when their previous contracts ended. Column 2 considers only US and Canadian nationals, whereas column 3 considers only players with foreign nationality. Column 4 considers only players in seasons 2011–2017, and column 5 only players in seasons 2007–2010. Across all these specifications, the main result on $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$ still remains qualitatively robust, with the estimate varying between -0.07 and -0.22. However, we would encourage significant caution in interpreting the set of results in Table 3. In these smaller and more restricted samples, the degree of MLS mobility represented decreases substantially compared with the main results. This would suggest that the estimates of $Corr_{it}(\hat{\alpha}_i, \hat{\phi}_{j(it)})$ are likely biased downward by limited mobility bias, as

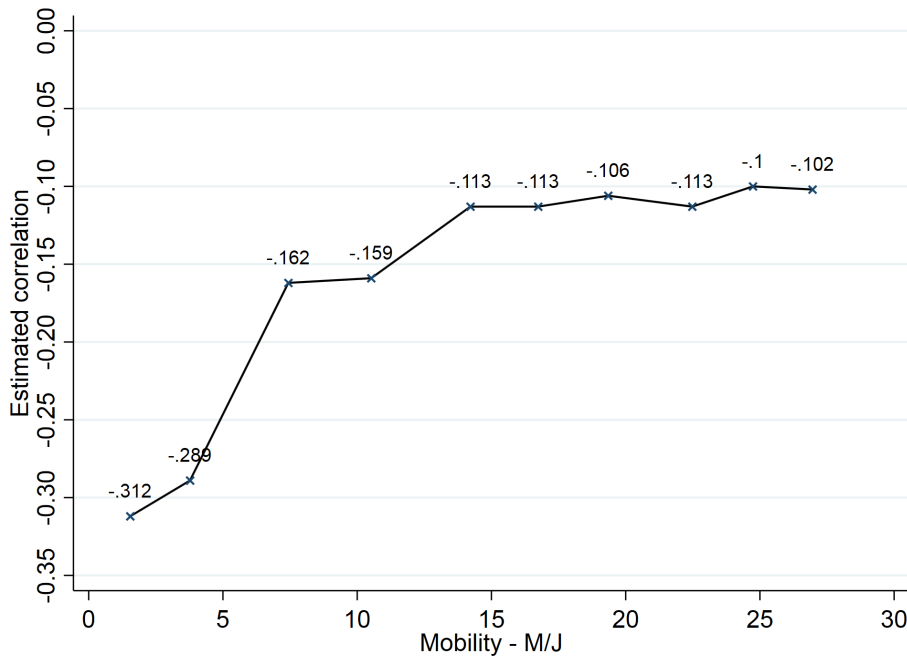


Figure 6. Limited Mobility Bias and the Estimated Correlation Between Player and Team Fixed Wage Effects, $Corr_{it}(\hat{\alpha}_i, \hat{\varphi}_{j(it)})$

Notes: Estimates are calculated by drawing random 10%, 20%, etc. samples of player-year observations from those who started at least one game during the year. *M/J* gives the average number of player moves per team in the resulting dataset’s largest connected set of players and teams.

demonstrated by Figure 6. For example, this bias is more likely to explain why the estimated correlation of player and team fixed effects is closer to zero for US and Canadian nationals than it is for foreign nationals, rather than there being any innate difference between these two groups in how high wage players are allocated to high wage teams in MLS.

Further Discussion

Taken together, the evidence above suggests that our estimates of α_i reflect individual player productivity but that our estimates of φ_j are not reflective of team-level productivity. Although the most productive players do play for the most productive teams, the most productive teams do not pay high wage premiums. One potential explanation for the estimated small negative correlation between player and team wage premiums in MLS is the existence of some non-pecuniary attraction or compensating wage differential that players receive for playing for high productivity and well-supported teams. In other words, due to their preferences, the best players might be willing to accept a lower team wage premium in order to play for the best teams. For example, players may perhaps receive extra utility from doing well in the League. Alternatively, playing for a more productive team may lead to greater outside or future earning potential, either through endorsements or corporate sponsorship deals or through transferring to another league.

However, most players in MLS generally have little free agency or bargaining power with the teams. In other words, once players enter MLS and are assigned to teams, those teams have significant monopsony power over them. This is reflected by the fact that the players union is prominent and has consistently pushed to increase the agency of its members during the annual collective bargaining agreements. For example, under pressure from the union, MLS introduced the reentry draft, first taking place before the 2011 season, which increased the possibility of movement within the League for veteran players. Further, this is not a labor market where wages are posted but where wages are negotiated after a player and a team have decided they are likely to be a good match. The fact that some teams are successful in MLS could be due to them being systematically more able or willing to apply their high monopsony power over players. This is especially plausible when factoring in the presence of the salary cap and other roster controls in MLS. If teams are generally more productive because they are tougher salary negotiators, then this explanation would also be consistent with the sports intuition that high productivity players in fact do play for high productivity teams, only that we are unable to identify

this from wage data alone in this particular labor market case. To offer some support for this conclusion, we consider the correlation between the players' own estimated values of α_i and the mean of their teammates, or coworkers, over all player-year observations in the sample, i.e., $Corr_{it}(\hat{\alpha}_i, \tilde{\alpha}_{J(it)})$, where $\tilde{\alpha}_{J(it)}$ is the mean level of α_i among the teammates of player i . As Lopes de Melo (2018) explained, this provides a measure of coworker segregation from the AKM wage model estimates. In MLS, we estimate this correlation to be 0.22. There is suggestive evidence of positive coworker sorting in this labor market. This implies that the distribution of player types or ability levels is not the same across firms, despite the finding that high wage players do not play for high wage teams, as evidenced by our small and negative estimate of $Corr_{it}(\hat{\alpha}_i, \tilde{\alpha}_{J(it)})$.

As mentioned before, Drut and Duhautois (2017) found very different results in Italian football, with evidence from wage data of substantial positive assortative matching between players and teams. Ribeiro and Lima (2019) also found that players in Portugal receive wage increases when they move to teams further up in the domestic leagues, indicating that the high productivity teams in Portugal pay high wages. However, professional Italian and Portuguese football and US soccer differ in one key respect, that being competitive balance. Italian football, like most European leagues, has a much clearer hierarchy, as evidenced by Drut and Duhautois (2017) finding that the team fixed effects can account for over a quarter of wage variation there. The equivalent estimate from MLS is just 1%, despite there being significant differences in these premiums between teams. The Portuguese first division is even less competitive: Only four different teams have won the league since its inception in 1934. Therefore, one plausible explanation for the differences between MLS and Italian or Portuguese football is that the presence of regulation and the salary cap has largely prevented the most productive teams from paying high wage premiums to attract the most productive players in the market for football talent. Instead, some teams plausibly use and work with that regulation more effectively than others, thus exercising significant monopsony power over their players, bargaining down wages, and being able to put talent on the pitch at lower average cost.

Conclusion

In this article, we show that high wage professional footballers do not play for teams that pay a high wage premium, at least so far as MLS is concerned, despite intuition and some existing empirical evidence that suggests the most talented, or productive, footballers do play for the most successful teams.

We show that the player fixed effects in an estimated AKM wage equation do seem to reflect individual player productivity. However, the team fixed effects from that equation do not appear to reflect overall team productivity: Teams that pay a high wage premium have tended to have less success in the League. There are plausible reasons why this may be the case in a football talent market, particularly one that is regulated like MLS, which are consistent with more productive teams being able to extract a greater share of the matching surplus upon acquiring a player. If this is the case, then our results may not generalize to other markets with fewer salary regulations, such as the sports talent markets outside of North America.

Our findings also provide a cautionary example by adding to a growing literature which suggests that wage data alone can be insufficient to identify the nature of assortative matching on productivity between firms and workers.

Endnotes

¹For example, the salary cap does not cover young “homegrown” players from the USA.

²For example, Francesco Totti played for his home-town team AS Roma for his whole career, despite reported huge offers from Spain for his services (<https://www.marca.com/en/football/real-madrid/2020/05/18/5ec2972422601db6648b45a3.html>).

³It does not include team- or individual-level performance payments. However, according to the League's annual bargaining agreements, team-level performance payments are a tiny fraction of overall salary payments. Individual performance payments that are considered achievable are included in the team's total salary figure, which is subject to a cap. We have compared total team salaries to the salary cap and found that reported total team salaries in fact exceed the cap, suggesting that individual performance payments are not large.

⁴The analysis covers 23 teams, since one was dissolved in 2014; 2007: 13, 2008: 14, 2009: 15, 2010: 16, 2011: 18, 2012–2014: 19, 2015–2016: 20, and 2017: 22.

⁵<https://www.mlssoccer.com/stats>. The incomplete nationality data from the MLS was cross-checked and supplemented using internet searches on players.

⁶See <https://www.mlssoccer.com/league/official-rules/mls-roster-rules-and-regulations> for details on the rules and regulations regarding how teams acquire talent and how players can move between teams.

⁷Johnson & Hall (2018) also found that wage premiums peak at age 27 for professional sportsmen in the US in the case of basketball players in the NBA.

⁸This is illustrated by Figure 1. In 2017, for example, a player at the 90th wage percentile earned over 11 times more than a player at the 10th percentile.

⁹In MLS, a win earns 3 points, a draw 1 point, and a defeat 0 points. Cumulative points determine final league standings and qualification for playoffs.

¹⁰See for example Forbes 2018 value estimates of MLS teams: <https://www.forbes.com/sites/chris-smith/2018/11/14/mls-most-valuable-teams-2018>

¹¹For example, in 2016, Forbes estimated that Kaká earned \$2 million, more than 25% of his salary, through endorsements (<https://www.forbes.com/pictures/mlh45fhdf/1-kaka-orlando-city-sc>).

References

- Abowd, J., Creecy, R. H., & Kramarz, F. (2002). *Computing person and firm effects using linked longitudinal employer-employee data*. Longitudinal Employer-Household Dynamics Technical Papers. Center for Economic Studies, U.S. Census Bureau.
- Abowd, J., Kramarz, F., & Margolis, D. (1999). High wage workers and high wage firms. *Econometrica*, 67(2), 251–334.
- Andrews, M., Gill, L., Schank, T., & Upward, R. (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters*, 117(3), 824–827.
- Andrews, M. J., Gill, L., Schank, T., & Upward, R. (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society Series A*, 171(3), 673–697.
- Bartolucci, C., Devicienti, F., & Monzon, I. (2018). Identifying sorting in practice. *American Economic Journal: Applied Economics*, 10(4), 408–38.
- Becker, G. (1973). A theory of marriage: Part I. *Journal of Political Economy*, 81(4), 813–46.
- Brownlee, E., & Lorgnier, N. (2017). A case study of Major League Soccer. In N. Chanavat, M. Desbordes, and N. Lorgnier (Eds.), *Routledge handbook of football marketing* (pp. 73–87). Routledge.
- Bryson, A., Rossi, G., & Simmons, R. (2014). The migrant wage premium in professional football: A superstar effect? *Kyklos*, 67(1), 12–28.
- Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1), S13–S70.
- Coates, D., Frick, B., & Jewell, T. (2016). Superstar salaries and soccer success: The impact of designated players in Major League Soccer. *Journal of Sports Economics*, 17(7), 716–735.
- Drut, B., & Duhautois, R. (2017). Assortative matching using soccer data. *Journal of Sports Economics*, 18(5), 431–447.
- Eeckhout, J., & Kircher, P. (2011). Identifying sorting: In theory. *Review of Economic Studies*, 78(3), 872–906.
- Filippin, A., & van Ours, J. C. (2015). Positive assortative matching: Evidence from sports data. *Industrial Relations: A Journal of Economy and Society*, 54(3), 401–421.
- Frick, B. (2007). The football players' labor market: Empirical evidence from the major European leagues. *Scottish Journal of Political Economy*, 54(3), 422–446.
- Gandelman, N. (2008). Mobility among employers and assortative matching. *Journal of Sports Economics*, 9(4), 351–370.
- 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.
- Johnson, C., & Hall, J. (2018). Do National Basketball Association players need higher salaries to play in high tax states? Evidence from free agents. *Applied Economics Letters*, 25(5), 359–361.
- Kuethe, T., & Motamed, M. (2010). Returns to stardom: Evidence from U.S. Major League Soccer. *Journal of Sports Economics*, 11(5), 567–579.
- Lawson, R. A., Sheehan, K., & Stephenson, E. F. (2008). Vend it like Beckham: David Beckham's effect on MLS ticket sales. *International Journal of Sport Finance*, 3(4), 189–195.
- Lopes de Melo, R. (2018). Firm wage differentials and labor market sorting: Reconciling theory and evidence. *Journal of Political Economy*, 126(1), 313–346.
- Lucifora, C., & Simmons, R. (2003). Superstar effects in sport: Evidence from Italian soccer. *Journal of Sports Economics*, 4(1), 35–55.
- Ribeiro, A. S., & Lima, F. (2019). Football players' career and wage profiles. *Applied Economics*, 51(1), 76–87.
- Shimer, R., & Smith, L. (2000). Assortative matching and search. *Econometrica*, 68(2), 343–370.
- Thrane, C. (2019). Performance and actual pay in Norwegian soccer. *Journal of Sports Economics*, 20(9), 1051–1065.
- Wooten, J. J. (2018). Rivalries in US professional soccer. *International Journal of Sport Finance*, 13(2), 105–118.