Universidad de SanAndrés

Departamento de Economía

Licenciatura en Economía

The inefficiencies of a billionaire market

Mateo Juan Criado

Legajo 29.042

mcriado@udesa.edu.ar

Supervised by: Jorge Alan Acosta

Spring 2020

Contents

List of Figures

List of Tables

Abstract

This paper investigates the labor market of the English Premier League by studying the relationship of wages with performance and popularity. The results show a weak link between performance and wages and a strong link between popularity and wages. Furthermore, when analyzing lags in salary adjustments to address whether the previous results had been a consequence of information asymmetries, the outcomes show that performance is now a worse indicator than before. The findings in this paper might imply that, in the English Premier League, clubs determine salaries considering the revenue that players generate and not only how well they play.

Keywords: sports, labor market, weighted output, market efficiency, salary adjustments, performance, popularity

1 Introduction

In the following sections, I will provide evidence that, if sports output is player performance, market efficiency is being violated in the labor market of the English Premier League. I first present several indices for player performance, which rely on compound algorithms to produce a unique rating. I start by analyzing whether the salary of players is a function of their performance. Then, based on the hypothesis that professional sports aim to maximize monetary benefits, and that this does not necessarily imply maximizing sports results, I consider whether player popularity –because it means, for example, more tickets and jersey sales and increased sponsoring – is an indicator of wages. Lastly, I explore lags in performance-based wage determination to address information asymmetry problems.

According to the modern neoclassical economics approach, in an efficient market with rational agents, workers should receive a salary that is directly related to their productivity. In particular, in football, according to this school, a player should receive compensation that is based on its input to help the team win, and this is why players who perform better should earn more. In simple words, if the football market was efficient, we should be able to model players' salaries as a function of their performance. Similarly, if the input in this market was a contribution to the financial benefits of the clubs –such as stadium assistance, clothing sales, and sponsoring – then we should be able to model salaries as a function of this contribution, which I will define as player popularity.

My work was motivated by Massey and Thaler (2013) and Adoumie (2019) papers, which studied the NFL draft pick and the NBA labor market, respectively. I then wished to replicate some of their studies in the football English Premier League and to widen further the investigation. In particular, I will explore two extensions: (i) Can player popularity explain salary?, and (ii) Is the stochastic nature of player performance, a problem addressed by negotiating salaries *ex post*? I will study the

English Premier League 2016/2017 season because of the availability of data and to be able to study possible lags in salary determination.

The main results show that player performance is not solidly linked to compensation, which means that the English Premier League is inefficient at assigning wages based on player performance. On the other hand, the results suggest that popularity is an effective indicator of salary determination. Finally, the results regarding lagged salary adjustments suggest asymmetric salary rigidity. These might have implications on various fields such as behavioral economics, information economics, and industrial organization.

The Annual Review of Football Finance in 2019 assembled by Deloitte[1] briefs that, during the 2017/2018 season, the revenue of the English Premier League was over €5.440 million and that the expected revenue for the 2019/2020 season was near €6 billion. If the latter occurs, it would mean an increase of 140% from the €2.5 billion in the season 2012/2013. Football is a rapidly growing business, and analyzing labor in such a big market –in terms of money flow, growth, and people employed – becomes increasingly important to understand the financial behavior of its agents, and therefore understanding the underlying factors of the market as a whole.

2 Wages and productivity in sports

In compliance with the propositions of the neoclassical economic theory regarding efficiency and rational expectations, wages should relate to productivity. In most industries, skills and human capital are difficult to measure. What makes the sports industry of excellence to study pay and productivity is that it is relatively easy to measure work performance, that is, productivity. 'Labour markets in professional sports' by Rosen and Sanderson finds that to the extent that the Coase theorem

[1]Multinational professional services network. It is one of the Big Four accounting organizations and the largest professional services network in the world by revenue and number of professionals.

applies to sports (*a priori* our case), the deadweight losses that relate to player allocation are eliminated. Besides, as it is also expressed in Rosen and Sanderson (2001), there are times where performance is not everything: "memories of past performances attract fans to see the aged stars, not necessarily their current productivity on the team". The latter quotation will be addressed in section 5.2.

The work of Carmichael et al. (2011) investigates the relationship between playing success and commercial success in team sports. They find evidence that "on-field success can be directly related to players' skills and abilities" and that "wage expenditure is also shown to systematically reflect player skills and performance". Even if the latter was true in the EPL, it would not necessarily convey that the market is efficient. However, their conclusions can be used as a starting point and motivate further detailed analysis of the football labor market.

'The loser's curse: Decision making and market efficiency in the National Football League draft' by Massey and Thaler (2013)[2], is one of the most well-known papers on the sports labor market and it studies the United States NFL draft market, where they compare the market value of draft picks with the surplus-value to teams provided by the drafted players. The paper demonstrates inefficiencies in the draft market –the league's young talents – and irrational economic actions of the decision-makers, which in turn were consistent with psychological research. Specifically, it concludes that "top draft picks are overvalued in a manner that is inconsistent with rational expectations and efficient markets and consistent with psychological research".

My work was inspired by Adoumie (2019), an honors thesis at the University of California, Berkeley. Adoumie worked on a similar investigation to Massey and Thaler (2013) but for the NBA. His work explored basketball and included not only the draft market of rookies but also the "free no-salary-cap" market of experienced players. Adoumie's paper finds that in the NBA "performance is a poor indicator of

[2]Richard Thaler is a behavioral economist who won the Economics Nobel Prize in 2017.

compensation", which keeps in line with Massey and Thaler (2013) conclusions for the NFL league. Adoumie adapted Massey and Thaler (2013) draft methodology to a no-salary-cap market. I will use the same adaptation of the methodology because the NBA experienced pool salary rules are wholly comparable to the ones in the English Premier League, except for the teams' upper limit for expenditure, which should not affect our analysis.

McHale et al. (2012) paper 'On the development of a soccer player performance rating system for the English Premier League' affirms that, although rating players in an individual sport is relatively easy and, likewise, rating teams in team sports is also easy, there is a particular difficulty in rating individually the players of team sports. In their work, they explain the development of the EA Sports Player Performance Index that aims to rate all players, regardless of their position or specialty, with a unique scoring system. The necessity of having overall performance indices rather than common statistics, motivated me to investigate some of the best available metrics to use them in this work. My paper will use this academy-developed index, apart from others, for measuring player performance.

Furthermore, I will explore whether salary determination can occur regardless of player performance. In particular, I will try to analyze if there is hiring occurring because of name or brand recognition. For this study, I will have a base in Rosen and Sanderson (2001), as commented before, and on Frick (2007) 'The football players' labor market: empirical evidence from the major European leagues', who, after studying the labor market of the major European Football Leagues, found that the volume of Google searches for a particular player, was (positively) related to its salary.

3 Metrics method

Although analyzing classical statistics like goals, assists, clean sheets, or distance covered could be interesting, these do not reflect an overall performance because they are sensitive to the player position or specialty. It is vital to use data that measure player performance as a whole, such as McHale et al. (2012) explains. For this paper, I decided to use all statistic-based measurements: the EA Sports Player Performance Index, developed by a collaboration of academia, media, and professional soccer leagues, the ICT Index, developed by Fantasy Premier League, WhoScored ratings, and the overall player rating of FIFA, an EA Sports football video-game. The first index responds to academia, the second to the official institutions of the English Premier League, media uses the third one, and entertainment uses the fourth one. Therefore, I will cover with very solidly developed and widely respected indices all the main participants of the football market. I will briefly explain them below.

3.1 EA Sports Player Performance Index

The EA PPI is a classification system for soccer players used in the two highest levels of soccer in England: The Premier League and the Championship. Its development was primarily carried out by academia and included the collaboration of professional soccer leagues and an association of media, and described in McHale et al. (2012) 'On the development of a soccer player performance rating system for the English Premier League'. A linear combination of 6 subindices makes up this index.

- Subindex 1 - Modeling match outcome: It can be thought of as a simpler version of the whole EA PPI Index. Explaining its methodology in a few words might only lead to further confusion because of the complexity it carries, but this graph from Klaiber (2016) can help:

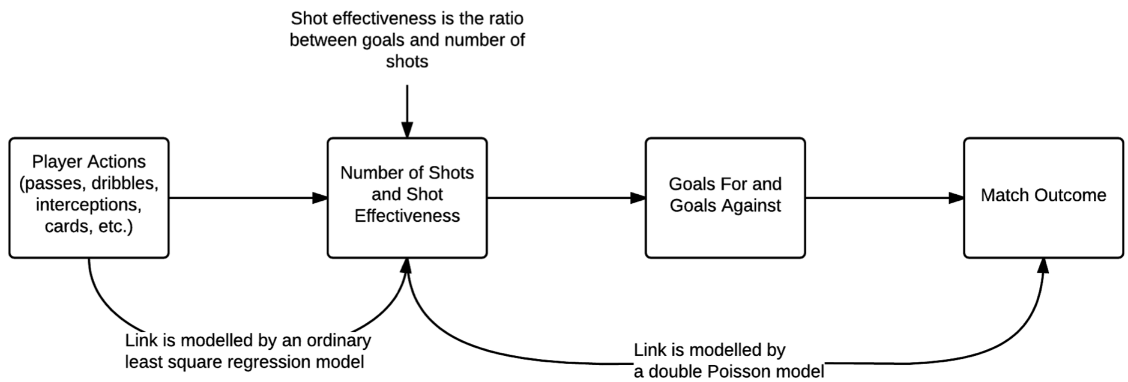

Figure 1: Klaiber's (2016) explanation of the process that relates player actions to the match outcome for EA PPI subindex 1.

We can summarize the first subindex as the "match contribution" of a player.

- Subindex 2 - Point Sharing Index: It grants points to players that played when their team won league points. It is calculated by multiplying the playing time ratio that the player was on the pitch over the total sum of minutes of all players, with the number of league points that the team won. It shares out the points won by the team among the players who participated in the game. We can summarize this subindex as a "winning performance".

- Subindex 3 - Appearance Index: It is similar to the second subindex, but it does not relate players to winning performance. Instead, it awards points merely for playing. This subindex divides the total amount of points of every team in the league among players, considering the number of minutes each one played. It can be summarized as "match appearances".

- Subindex 4 - Goal-scoring Index: The points in this index are the number of goals of a player multiplied by the average number of league points per goal (total number of league points divided by the total number of goals). We can simplify it as "goals scored".

- Subindex 5 - Assists Index: The analogous to the Goal-Scoring Index but for assists. An assist is credited to a player for passing or crossing the ball to the scorer.

- Subindex 6 - Clean-Sheets Index: It rewards players for not receiving goals. To keep a balance of the overall index, they take that total points awarded for assists equal total points for clean sheets. According to the data they studied, on the total defensive actions, goalkeepers perform 21%, defenders 13% each, midfielders 5% each, and strikers 3% each. Thus, if a clean sheet awarded 3 points, a goalkeeper should receive for subindex 6 the amount of $0.21 * 3 = 0.63$.

Finally, the final index is the weighted sum of the points achieved on each subindex, which they define as:

$$I = 100 \times (0.25I_1 + 0.375I_2 + 0.125I_3 + 0.125I_4 \\ +0.0625I_5 + 0.0625I_6)$$

3.2 ICT Index

The ICT Index is a football statistical index developed by the Fantasy Premier League to assess a player as an asset. The index generates a score for each key field of the ICT, which are: Influence, Creativity, and Threat. The ICT index condenses more than 40 match event statistics and is generated by the weighted sum of the following sub-indices:

- Influence: assesses the degree of impact of a player in the game, taking into consideration actions and events that may affect directly or indirectly the match outcome.

- Creativity: evaluates the performance in terms of generation of goal opportunities, but not only assists. While it analyzes the frequency of passes and crossing, it also considers pitch location and quality of the final ball. Moreover, it identifies the players most likely to supply assists.

- Threat: examines a player's menace on goal, so it rewards the ones most likely to score a goal. Attempts are the central action of this subindex, but it also looks at pitch location and "gives greater weight to actions that are regarded as the best openings to register a goal".

The final ICT Index is calculated as:

$$ICT = \frac{I + C + T}{10}$$

3.3 FIFA Ratings

The compilation of 300 fields and 35 attribute categories for each player build this index, and over 9000 people, including scouts, coaches, and fans, participate in the process. The amount of data used exceeds 5.4 million observations, among the approximately $18,000$ players and 700 clubs. The rating value of the players in this index ranges from 0 to 100.

Although the exact algorithm was not revealed by EA Sports[3], SOFIFA, a widely known website specialized in the EA FIFA game, offers a calculator of FIFA player rating, which boats high accuracy. The algorithm for calculation varies for each position and was revealed by Kevin Healey at Kaggle[4]. According to the evidence in Healey (2017), these revelations are impressively accurate. The estimations of the algorithms used by FIFA are the following:

1. Goalkeepers

$$GK = 0.21gd + 0.21gh + 0.21gp + 0.21gr + 0.11re + 0.5gk$$

[3]Division of Electronic Arts Inc. and owner of the FIFA video-game.

[4]Kaggle is a subsidiary of Google LLC that, according to Wikipedia, is an "online community of data scientists and machine learning practitioners"

2. Wing backs

$$RWB\&LWB = 0.04ac + 0.06sp + 0.1st + 0.08re + 0.12in + 0.08bl$$
$$+0.12cr + 0.04dr + 0.1sh + 0.07ma + 0.08sa + 0.11sl$$

3. Backs

$$RB\&LB = 0.05ac + 0.07sp + 0.08st + 0.08re + 0.12in + 0.07bl$$
$$+0.09cr + 0.04he + 0.07sh + 0.08ma + sa0.11 + sl0.14$$

4. Center backs

$$RCB\&LCB = 0.02sp + 0.03ju + 0.1sr + 0.05re + 0.07ar + 0.13in$$
$$+0.04bl + 0.1he + 0.05sh + 0.14ma + 0.17sa + 0.1sl$$

5. Defensive midfielders

$$RDM\&LDM = 0.06st + 0.04sr + 0.07re + 0.05ar + 0.14in + 0.04vi$$
$$+0.1bl + 0.1lo + 0.14sh + 0.09ma + 0.12sa + 0.05sl$$

6. Regular midfielders

$$RM\&LM = 0.07ac + 0.06sp + 0.05st + 0.07re + 0.08po + 0.07vi$$
$$+0.13bl + 0.1cr + 0.15dr + 0.06fi + 0.05lo + 0.11sh$$

7. Center midfielders

$$CM = 0.06st + 0.08re + 0.05in + 0.06po + 0.13vi + 0.14bl + 0.07dr$$
$$+0.02fi + 0.13lo + 0.17sh + 0.04ln + 0.05sa$$

8. Attacking midfielders

$$AM = 0.04ac + 0.03sp + 0.03ag + 0.07re + 0.09po + 0.14vi + 0.15bl$$
$$+0.13dr + 0.07fi + 0.04lo + 0.16sh + 0.05ln$$

9. Forwards

$$F = 0.05ac + 0.05sp + 0.09re + 0.13po + 0.08vi + 0.15bl + 0.14dr$$
$$+0.11fi + 0.02he + 0.09sh + 0.05so + 0.04ln$$

10. Wings

$$W = 0.07ac + 0.06sp + 0.03ag + 0.07re + 0.09po + 0.06vi + 0.14bl$$
$$+0.09cr + 0.16dr + 0.1fi + 0.09sh + 0.04ln$$

11. Strikers

$$S = 0.04ac + 0.05sp + 0.05sr + 0.08re + 0.13po + 0.1bl + 0.07dr$$
$$+0.18fi + 0.1he + 0.05sh + 0.1so + 0.03ln + 0.02vo$$

The attribute names for each abbreviation used in the formulae are displayed below in Table 1:

Table 1: Attribute names by the abbreviation

Attribute	Abbreviation	Attribute	Abbreviation
crossing	cr	stamina	st
finishing	fi	strength	sr
heading accuracy	he	long shots	ln
short passing	sh	aggression	ar
volleys	vo	interceptions	in
dribbling	dr	positioning	po
curve	cu	vision	vi
free kick accuracy	fr	penalties	pe
long passing	lo	composure	cm
ball control	bl	marking	ma
acceleration	ac	standing tackle	sa
sprint speed	sp	sliding tackle	sl
agility	ag	GK diving	gd
reactions	re	GK handling	gh
balance	ba	GK kicking	gk
shot power	so	GK positioning	gp
jumping	ju	GK reflexes	gr

Source: Own elaboration based on Healey (2017).

Unfortunately, we do not know the methodology that EA Sports uses to determine the value for each attribute in FIFA. Nonetheless, for the purpose of the analysis, because of the positive international reputation of EA Sports and the many years they have been in business, I assume and believe that the more than 9000 people who participate in the process make the ratings non-biased and robust.

3.4 WhoScored Ratings

WhoScored ratings are popularly considered accurate and respected, and the world of football uses them as performance indicators. They are currently used by media giants such as The Guardian, ESPN, AS, football clubs, and bookmakers. Although the exact algorithm is not revealed, WhoScored explains that their ratings are exclusively statistical, unique, and calculated during the game, with data provided by the reliable sports data-giant Opta. WhoScored affirms that there are “over 200 raw statistics" included in the calculation of a player/team rating, weighted according to their influence within the game. "Every event of importance is taken into account, with a positive or negative effect on ratings weighted concerning its area on the pitch and its outcome" they claim. Their ratings go from zero to ten, with all players regardless of position and club starting at a rating of six.

4 Data

This paper analyzes the 2016 − 2017 season. The reasons for this decision are both the willingness to study probable lags in salary determination and the availability of data –particularly because of the EA PPI Index –, whose information for the last two seasons is not available.

In the dataset, information for performance was gathered from several sources such as FPL[5], WhoScored, SOFIFA, Kaggle. All the sources are first handed and widely reliable. Furthermore, I randomly cross-checked observations from several variables with databases from other sources that provided the same data to make sure the information was precise. The players considered for the analysis were those that, after deleting the 0-minutes-played players, played the average number of minutes or more. I did this because it ensures that the player performance is being measured with a reasonable amount of data, making the data-set robust to outliers. The strategy is to mitigate players that might have played exceptionally well or exceptionally bad over a few games and would therefore not represent the reality I wish to analyze.

EA PPI is the index with most restricted data: among the 247 players that meet the requirements, there was data availability for 189 which, although not desirable in the first place, is a good portion indeed. All of the indices, provide, with their specific methodology, an overall rating of the players' performance; nonetheless, we

[5] Acronym for Fantasy Premier League

must know that measuring performance might be far more complex and that the perfect recipe for computing it might have not been discovered yet.

The popularity data was retrieved from Google Keyword Planner and Google Trends in the form of a volume of searches. The former provides the average monthly volume of searches for key terms, which in this case are the names of the players. However, Google Keyword Planner provides average search volume for the exact term, that is, the key-term "Sergio Agüero" does not provide data for key-terms "Kun Agüero" or "Agüero". To overcome this hurdle and make the popularity measured in monthly searches more realistic, the key-terms in the data-set used to retrieve the information includes both full name and only last name. Moreover, for players that are well-known only for their first names or aliases, for instance, "Willian" or "Fernandinho", I retrieved their data specifically. Furthermore, for players whose last names were common or repeated within the database, either as First or Last names, for example, "James", and for players whose last names are widely known for anything else rather than the football player itself, for example, "Philips", I only considered the searches volume for their full name. Finally, for top searched players, and with the sole purpose of improving the quality of the data, the search-volume figures were studied deeply –and corrected when necessary – with the Google Trends tool, that offers not volume measured in absolute numbers but, in turn, relative popularity. For each season, the popularity data corresponds to the period comprised between June 1st to May 31st of the following year, for instance, June 2016 - May 2017. The latter is due to the nature of the English Premier League that starts in June and finishes in May.

Finally, I retrieved player compensations from Spotrac, one of the largest online "sports team and player contract" resources on the internet. These compensations do not include performance or objective-based bonuses, but solely basic salary. Since football player contracts are not available in public or open official databases, I cannot assure that players earn the exact figure that the database provides. However, after randomly cross-checking salaries in several online databases, well-known media, and EA FIFA estimations, I consider that the information provided by Spotrac is more than acceptably realistic and will be useful for my analysis.

Below, Table 2 helps us to have an overview of the variables that are to be studied.

Table 2: Descriptive statistics of the dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Salary 16-17	247	3,275k	2,350k	228k	1,666k	4,160k	15,600k
Salary 17-18	202	2,930k	2,349k	312k	1,820k	4,420k	15,600k
Minutes	247	2,398	580	1,439	1,913	2,899	3,420
EA PPI	189	357	185	18	230	476	1,053
ICT	247	131	75	26.8	78.8	153	453
WhoScored	247	6.92	0.295	6.38	6.68	7.11	7.81
FIFA OVA	247	78.6	4.36	68	75	82	90
Searches	247	46,755	68,822	380	8,300	50,800	450k

As previously stated, all of the variables correspond to the season 2016-2017, except for the indicated salary.

5 Results

5.1 Performance and wages

Considering that the salary determination occurs in a free no-cap market, I will create a structure based on Massey and Thaler (2013) and Adoumie (2019) that will allow me to effectively scale the results for a better understanding of the analysis. More specifically, the process consists of assigning the value 1.0 to the player with the highest salary and scaling all of the other wages according to this one: the best paid will have a salary of scale 1, and all the other players will have a scaled salary smaller than 1. Furthermore, I will create a salary ranking: the position in this ranking indicates that the player in the first position earns the highest salary, the player in the second position earns the second-highest salary, and this scheme is valid to the whole list.

Although the chosen indices measure the performance of the players each match, I will use the overall ratings for the whole season, which lasts 38 games in one year. Following the system done with the salaries, for each index, I will assign the best rating, a value of 1.0, and scale all of the other ratings to the best performer.

Since the purpose of the analysis is to study whether salaries are good indicators of performance, that is, if we can model remuneration as a function of performance, the best way to explore it is to weight the scaled performance ratings by the scaled salaries for each player. The expectations for an efficient market are that every player earns relatively the same or, in other words, that each player receives the same compensation for each marginal unit of performance. I display the results in the scatterplots below:

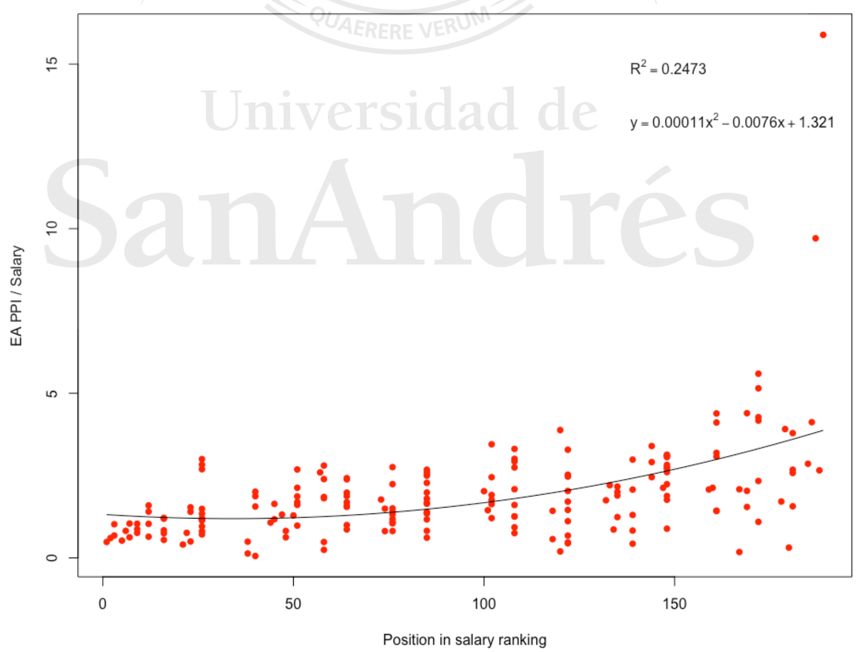

Figure 2: EA PPI weighted by salary to the position in salary ranking.

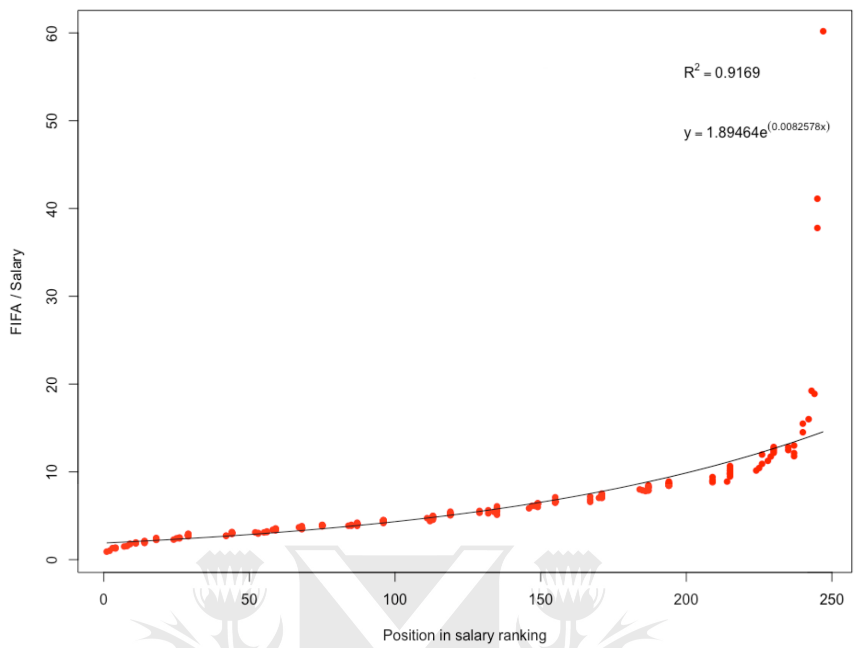

Figure 3: FIFA rating weighted by salary to the position in salary ranking.

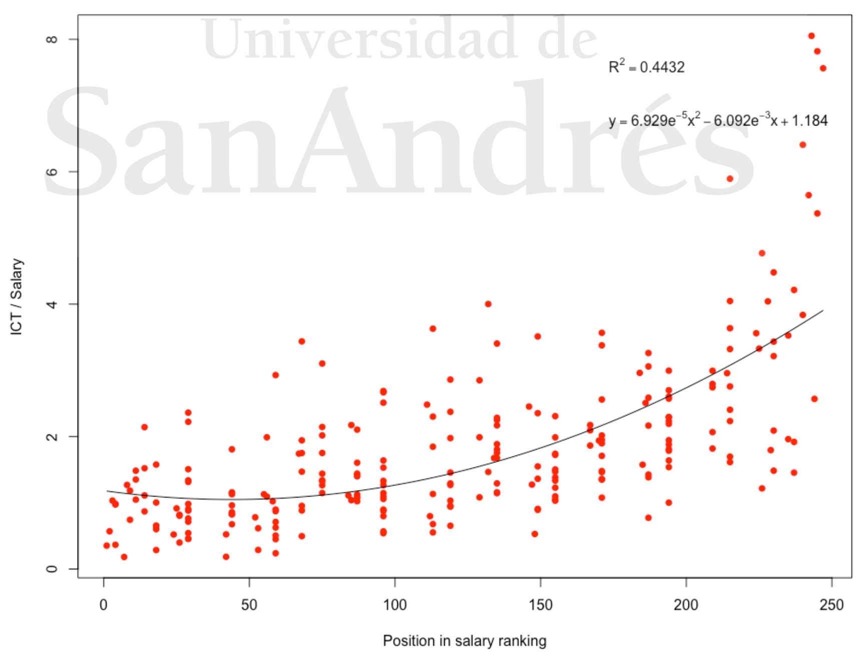

Figure 4: ICT Index weighted by salary to the position in salary ranking.

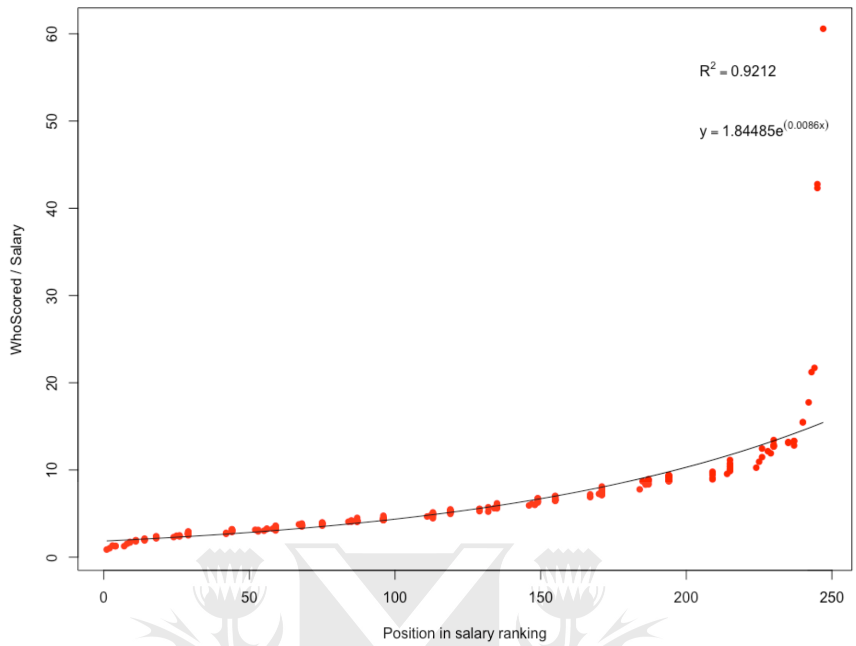

Figure 5: WhoScored ratings weighted by salary to the position in salary ranking.

In an efficient market, what we should have seen in the graphics above are trend-lines with a slope equal to zero: an utterly horizontal line. The results displayed in the figures above indicate that the English Premier League labor market is inefficient. We can appreciate the correlation between the weighted rating and the position in the salary rank, with two very high R^2 and quadratic and exponential fit-lines. The positive slope of the curve not only indicates that the players do not receive the same relative compensation, but that there is also a clear tendency regardless of the index used, where players with lower salaries have considerably higher weighted performance ratings. In simple words, players with lower retributions earn less money for each additional unit of performance. These results suggest that performance and salary are not linked in the way it is expected in an efficient market.

What I have shown before can be understood easily by thinking of an extreme situation: with these weighted ratings, should all the players perform identically, today's lower-paid players would still earn less money than today's best-paid players, *ceteris* *paribus*. Fundamentally, this means that the best performers are more productive for themselves and less productive for the clubs than the worse performers. We can see that the positions in the salary ranking can explain up to 92.12% of the variations in the data set when the expectations for an efficient market are that they should explain zero.

5.2 Popularity and wages

The results of the previous section indicate that if players with a different position in the salary ranking performed the same, one of them would receive more money than the other, *ceteris paribus*. Those results open interrogations for the causes of this inefficiencies and, to explore them, the primary hypothesis is that the club executives do not set salaries by basing their decisions on performance alone.

In this section, I will explore if popularity is a good indicator for compensation. This century is popularly known as the Information Age, and the intrusion of technology in our lives during the last two decades, particularly during recent years, has been immense and increasing without stopping. It should not surprise anyone that statistics prove this intrusion: according to Statista[6], during 2016 and 2017, the percentage of individuals using the internet in the United Kingdom was over 94%. Given this introduction to set a context and to substantiate my decision, I found appropriate to measure player popularity as the average monthly volume of searches on the internet.

The process that follows is analogous to the one described for the performance metrics: the volume of internet searches for each player is also scaled and weighted by the scaled salary. The results for this are displayed below:

[6]Statista is a widely known Germany-based portal for statistics that delivers more than one million statistics for over eighty-thousand topics, retrieving data from 22,500 sources.

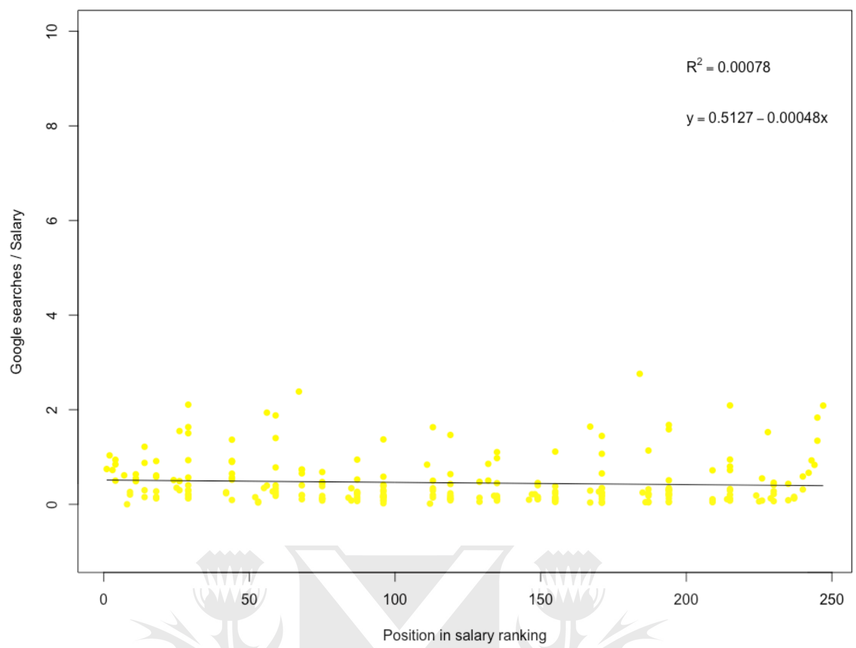

Figure 6: Searches volume weighted by salary to the position in salary ranking.

The outcomes here are surprising. The fit-line is quasi-horizontal [7], which means that the weighted searches to salary are virtually identical, almost constant, for every player. They suggest that popularity is a good indicator of compensation. The perfect efficiency would imply an R^2 equal to zero, meaning that the position of the player in the salary ranking has nothing to do with the weighted volume and that the latter in constant for every player.

5.3 Interpreting further

For visualization purposes, I made a correlation plot that displays Pearson's correlation between the weighted by salary indices and the position in the salary ranking:

[7] Please note that the scale in this plot is significantly smaller than the ones seen in previous figures.

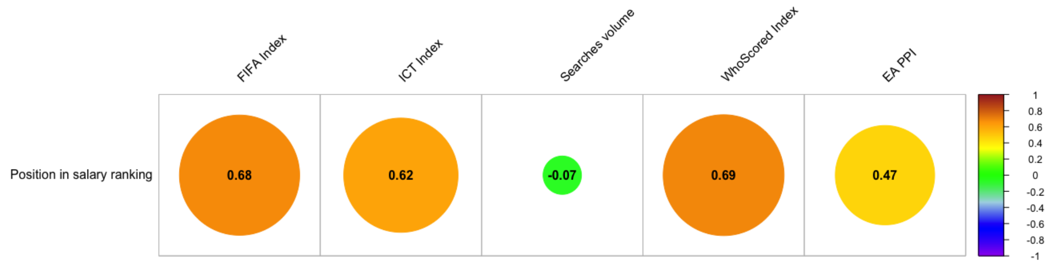

Figure 7: Correlogram for weighted metrics to the player position in salary ranking.

In an ideal market, we would expect these correlations to be zero. Once again, zero correlation would mean that the weighted performance, or that the weighted popularity, has nothing to do with the position in the salary ranking or, in other words, that each player earns what it deserves according exclusively to its staging or recognition, whichever case is. Endorsing the results shown in previous subsections, the top indicator is, once more, popularity.

Additionally, I decided to analyze what happened if I augmented the correlation between players of the same club by making the explanatory variable be the market-value of the team and not its individual position in the salary ranking of the whole league. Why? Each institution has its own rules and hiring policies, has its funding, and has its managers. Because of all these correlations among players of the same team, it appeared reasonable to perform this study to explore if there were any differences in the results. The market-value of the teams was retrieved from Transfermarkt[8]. I then run regressions of the weighted performances of each player against the market value of the squad to which they belonged. The results are displayed below:

[8]One of the largest and most well-known football-specialized websites in the world, which was founded in Germany.

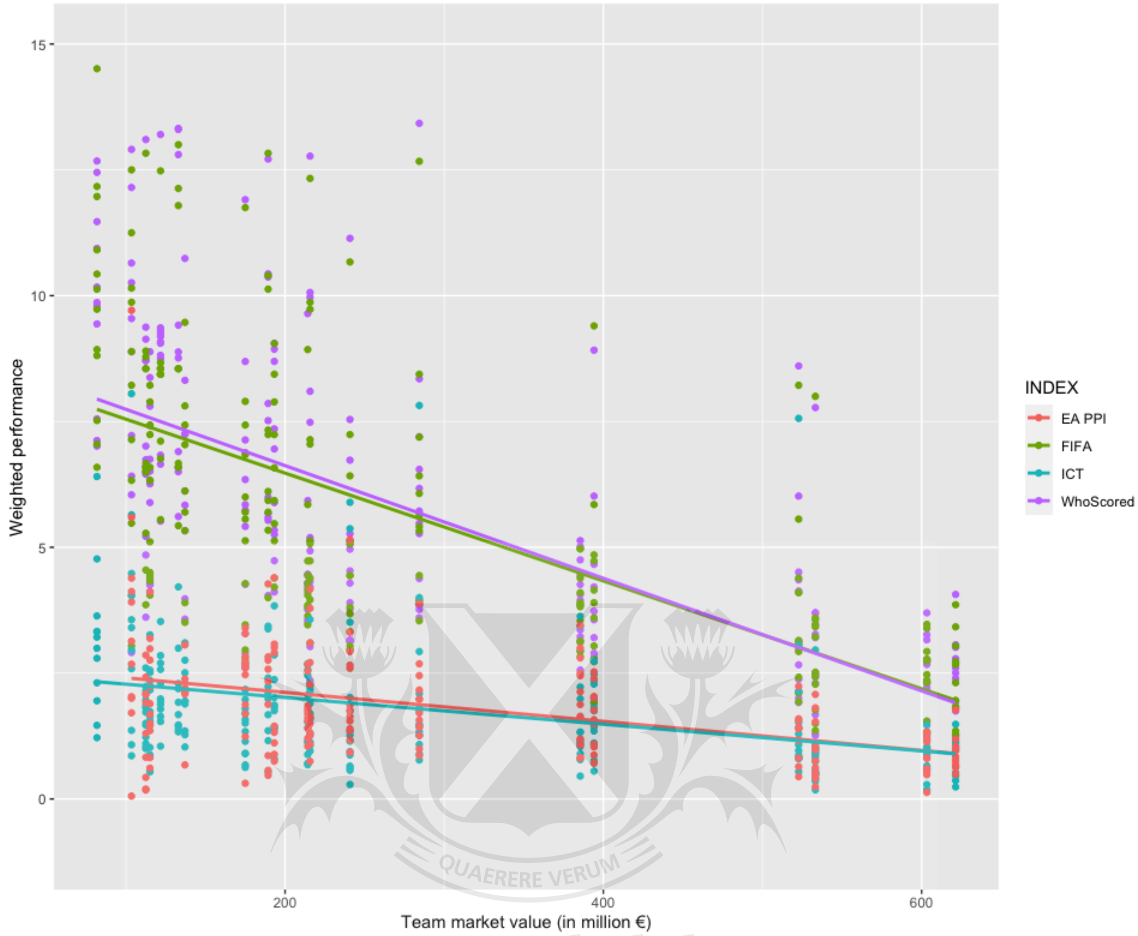

Figure 8: Weighted performances against team value.

These results are consistent with the findings in section 5.1. In particular, we can appreciate that those players who earn the most relative to their productivity belong to the teams with higher total market value, the richest, thus the ones who pay the highest salaries. It was somehow expectable to see these outcomes since it is probably true that the best players belong to the richest clubs. On that account, it was an indirect, or at least a highly-related way of visualizing what we have seen in the previous section, where we obtained that the players with lower salaries were the ones who performed the best in relative terms.

5.4 Testing for lags

In this section, I will explore whether salaries are elastic to player performance but with lags in the adjustments of contracts. To study it, I will make an analysis similar to the ones done in previous sections. In this opportunity, we want to know if salaries are adjusted after the season finishes since it is at this point where the results and the information of performance measurements are available for the executives of the clubs to make decisions regarding salaries.

The method to explore whether this is the case is by doing an analogous analysis but testing performance in t with salaries in $t+1$. In simple words, I will use the performance metrics results from EPL season $2016-2017$ and the salaries from EPL season $2017-2018$. Although there are naturally some players that did not stay in the league, the vast majority –over 80% – did; hence, the study is as well representative for elucidating any changes in the trends. The figures displayed below show the outcomes:

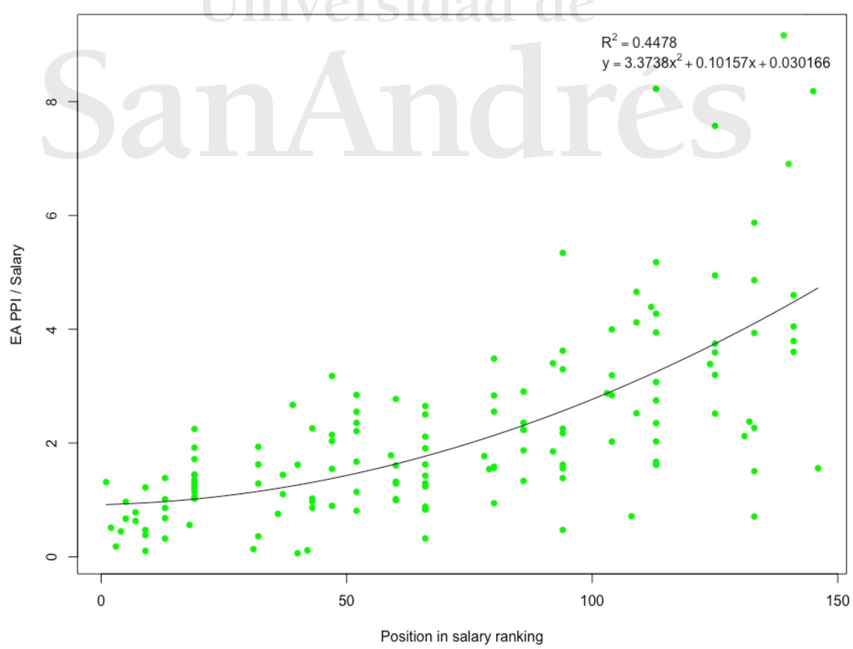

Figure 9: EA PPI weighted by salary to the position in salary ranking (salaries 2017-2018).

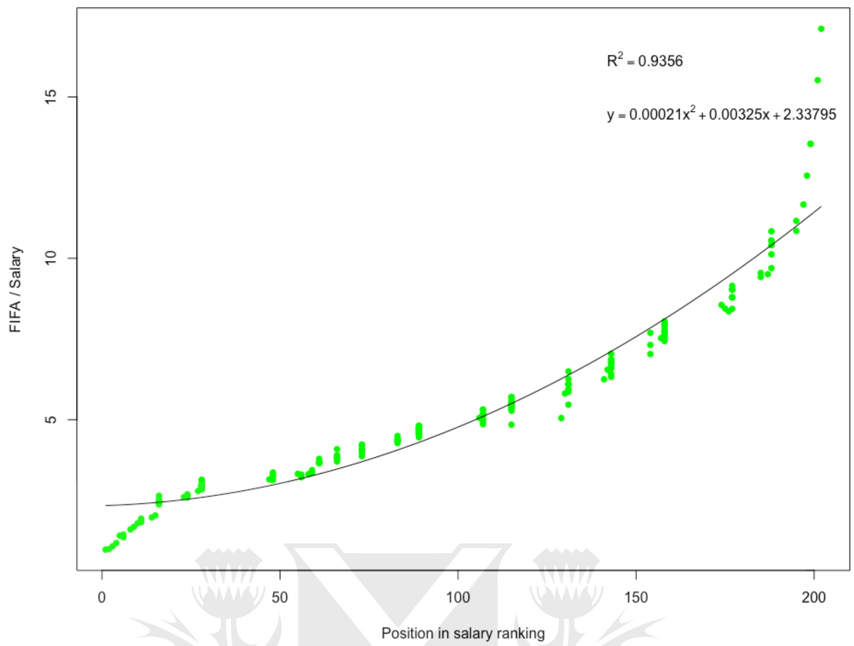

Figure 10: FIFA rating weighted by salary to the position in salary ranking (salaries 2017-2018).

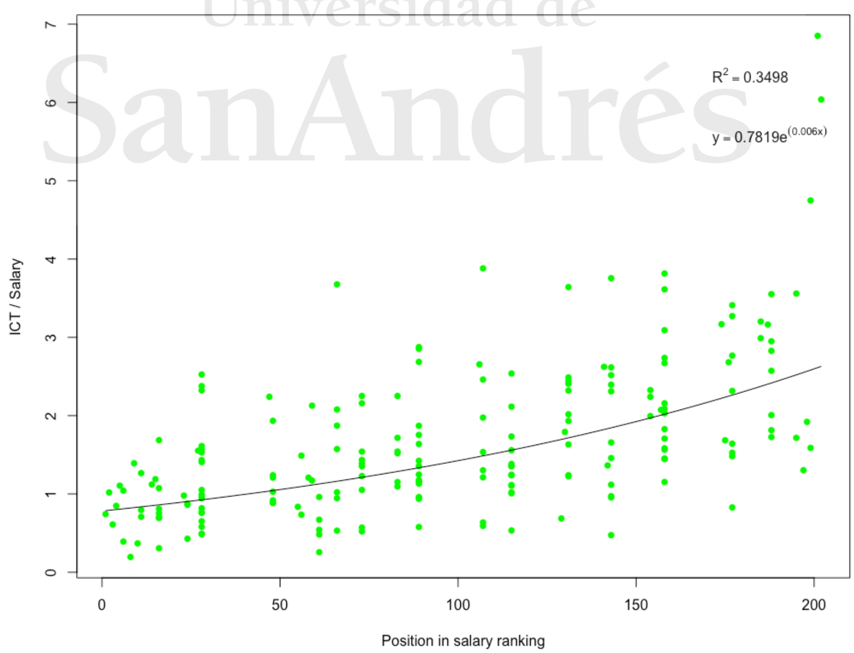

Figure 11: ICT Index weighted by salary to the position in salary ranking (salaries 2017-2018).

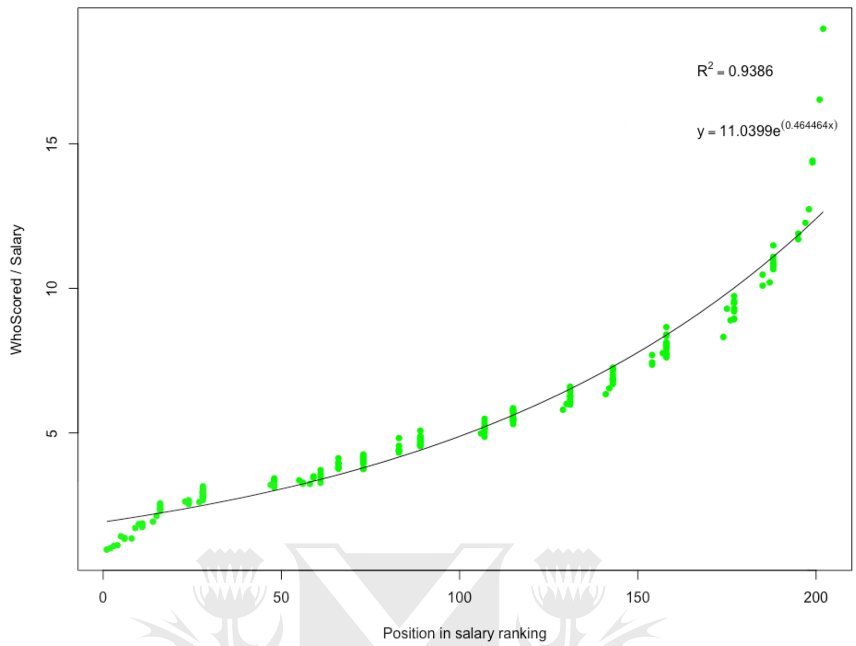

Figure 12: WhoScored ratings weighted by salary to the position in salary ranking (salaries 2017-2018).

The results displayed above show that the inefficiencies are greater in this case than when salaries are analyzed in the current season. This can be easily noticed by looking at the R^2 in each figure. The rise in inefficiency occurs in every case, except in the ICT Index, where the results show the opposite. Interpretations are left for further academic research.

6 Discussion and concluding remarks

Throughout the paper, I assessed efficiency in the labor market of the English Premier League from various perspectives. The presented results show that this market is inefficient if analyzing performance-based salary, supporting the conclusions made by Massey and Thaler (2013) for the NFL and by Adoumie (2019) for the NBA. On the other hand, when analyzing popularity-based salaries, the results suggest an almost perfectly efficient labor-market.

The possible understandings of the results are various. As stated previously, the measurement of individual performance in team sports is very complex, so the indices and ratings that exist are probably imperfect to explain reality, and consequently, this might be one of the causes of the ambiguities found. Another possible interpretation is that, since football is not an amateur sport where the output of the teams are the results of the games, the player compensations might not rely on the game results but, instead, on the multiple ways of revenue generation. Revenue of the clubs comes from many sources, such as TV rights, jerseys sales, ticket sales, marketing, transfers, sponsoring, and more. There is also a possibility, impossible to ignore, that there are teams that hire players because of their name or brand that generates revenue in other ways rather than playing.

Moreover, although difficult to quantify, it is popularly known that there are players who can contribute to group chemistry or morale, who can motivate, act as captains, or may teach the younger ones, albeit they might not be great performers. All of these reasons might also determine compensations.

When discussing popularity, I was impressed when I saw the results. Although I expected to see compelling outcomes because of the implications of popularity in revenue generation, I certainly did not expect these results to be remarkably different from the performance ones. However, we must not fall in love with them. On the one hand, they definitely tell us that popularity is a quality barometer for salary determination and they get along with the conclusions of Frick (2007). On the other hand, intuition tells us that popularity is mainly a result of the performance on-field. Hence, the results might also suggest that either the used performance indices might be inadequate indicators of real or human-perceived performance or that the top-performers are compensated by their popularity while the worse-performers not, or both. If there existed a minimum popularity threshold for which clubs remunerated popularity and also if performance explained popularity, it would elucidate why we saw efficiency in this case and not in the analysis of performance-based wages. However, I leave the latter for further academic research.

When speaking about the results of lag-testing in salary adjustments, the outcomes suggest that the inefficiencies were more substantial in this case. The implications for these might also be various, and some possible explanations –or a combination of them – in the order my intuition tells me, could be:

(i) Salaries are elastic to proficient performances yet inelastic to faulty ones. In other words, it is easier to raise or maintain wages than to reduce them. The asymmetry in the adjustments might explain the increase in inefficiency.

(ii) Executives determine salaries in such a way to make the top performers be happy in the club and make them feel important or super-stars to keep them in the team.

(iii) At the time of transfers, executives tempt with high salaries players who then do not repay on-field.

It is prime to understand the limitations of the data used, not only regarding the methodology of metrics but also the size of the sample, to appreciate the extent of the findings. Although the results bilaterally support and are supported by other studies in different sports, this paper studies one football league among the hundreds existing worldwide, that might have disparate rules, thus work differently. The latter might be the case, for example, of the MLS in the USA. Still, its relevance in the world of football is small, while the English Premier League is widely considered the most prestigious and high-level league, and also the oldest.

Concerning future works, it could be fascinating to make focus on contracts of "idols" or players that are emotionally represented by a particular team. Clubs might hire these players as a recognition for all they had performed in previous opportunities, or even for the sake of making the fans happier through a populist[9] decision. Further investigations might also make focus on studying football-markets dominated by few players and coach agents, like the cases of the Argentinian Superliga and the Mexican Liga MX. Also, they could study those where there are super-powerful club owners who control more than one club within the same league as the case of the Russian Premier League. Finally, studying the elasticity of contracts in the football market can be interesting. Players many times sign contracts for several seasons and if they perform adversely is difficult or impossible for club executives to trigger a drop in their salary. This difficulty exists *de facto* because almost every club agrees with the players the bonuses in contracts based on accomplishments of objectives but does not agree on salary reduction if players underperform, in part possibly caused because of the lack of credible and objective measurements of performance.

All the conclusions in this analysis have validity for football leagues where clubs operate as businesses, and results might vary significantly in countries where clubs are non-lucrative. The main findings in this paper include (i) inefficiency of the English Premier League labor-market regarding salary and performance and (ii) almost perfect efficiency regarding salary and popularity. If intuition is right and performance is the foremost indicator of output in sports, then this market does not work as the economic theory suggests. Finally, since all the players within the same team are paid out of the same funding, it should always be expected to see that the players in the winning teams earn more than the players in the worse-performing teams regardless of their personal performance. The latter is one of the reasons why this study alone cannot indicate that executives are managing the clubs irrationally.

[9] Derivation of populism in political science.

References

Adoumie, J. (2019). Violations of Neoclassical Economic Theory in the NBA Labor Market. *Economics Honors theses: University of California, Berkeley.*

Carmichael, F., McHale, I. & Thomas, D. (2011). Maintaining market position: Team performance, revenue and wage expenditure in the english premier league. *Bulletin of Economic Research*, *63*(4), 464–497.

Deloitte. (2013). 21% growth in commercial revenue helps premier league club revenue top £2.5bn for the first time. *Sport Business Group*. https://www2.deloitte.com

Deloitte. (2019). Annual Review of Fotball Finance. *Sport Business Group*. https://www2.deloitte.com

Frick, B. (2007). The football players' labor market: Empirical evidence from the major european leagues. *Scottish Journal of Political Economy*, *54*(3), 422–446.

Gerrard, B. (2010). Analysing sporting efficiency using standardised win cost: Evidence from the fa premier league, 1995–2007. *International Journal of Sports Science & Coaching*, *5*(1), 13–35.

Healey, K. (2017). Evaluating the sofifa.com rating calculator. *Kaggle*.

Klaiber, J. (2016). Soccer Player Performance Rating Systems for the German Bundesliga. *Universiteit Gent*.

Massey, C. & Thaler, R. H. (2013). The loser's curse: Decision making and market efficiency in the national football league draft. *Management Science*, *59*(7), 1479–1495.

McHale, I. G., Scarf, P. A. & Folker, D. E. (2012). On the development of a soccer player performance rating system for the english premier league. *Interfaces*, *42*(4), 339–351.

Rosen, S. & Sanderson, A. (2001). Labour markets in professional sports. *The economic journal*, *111*(469), F47–F68.

Appendices

A Performance and wages tables

Table 3: Weighted EA PPI fit-line regression

	Dependent variable:
	Weighted EA PPI
Salary Rank	-0.008
	(0.007)
Salary Rank (squared)	0.0001^{***}
	(0.00004)
Constant	1.321^{***}
	(0.288)
N	189
R^2	0.255
Adjusted R^2	0.247
Residual Std. Error	1.351 (df = 186)
F Statistic	31.891^{***} (df = 2; 186)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 4: Weighted FIFA fit-line regression

	Dependent variable:
	log(Weighted FIFA)
Salary Rank	0.008***
	(0.0002)
Constant	0.639***
	(0.022)
N	247
R^2	0.917
Adjusted R^2	0.917
Residual Std. Error	0.177 (df = 245)
F Statistic	2,716.514*** (df = 1; 245)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 5: Weighted ICT fit-line regression

	Dependent variable:
	Weighted ICT
Salary Rank	−0.006*
	(0.003)
Salary Rank (squared)	0.0001***
	(0.00001)
Constant	1.184***
	(0.172)
N	247
R^2	0.448
Adjusted R^2	0.443
Residual Std. Error	0.929 (df = 244)
F Statistic	98.890*** (df = 2; 244)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 6: Weighted WhoScored fit-line regression

	Dependent variable:
	Weighted WhoScored
Salary Rank	0.009***
	(0.0002)
Constant	0.612***
	(0.022)
N	247
R^2	0.921
Adjusted R^2	0.921
Residual Std. Error	0.179 (df = 245)
F Statistic	2,875.832*** (df = 1; 245)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

B Popularity and wages table

Table 7: Weighted searches fit-line regression

	Dependent variable:
	Weighted searches
Salary Rank	−0.0005
	(0.0004)
Constant	0.513***
	(0.062)
Observations	247
R^2	0.005
Adjusted R^2	0.001
Residual Std. Error	0.491 (df = 245)
F Statistic	1.193 (df = 1; 245)
Note:	*p<0.1; **p<0.05; ***p<0.01

C Lag-testing tables

Table 8: Weighted EA PPI fit-line regression for lags

	Dependent variable:
	Weighted EA PPI
Salary Rank	0.002
	(0.009)
Salary Rank (squared)	0.0002**
	(0.0001)
Constant	0.917***
	(0.293)
N	146
R^2	0.455
Adjusted R^2	0.448
Residual Std. Error	1.227 (df = 143)
F Statistic	59.784*** (df = 2; 143)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 9: Weighted FIFA fit-line regression for lags

	Dependent variable:
	Weighted FIFA
Salary Rank	0.003
	(0.003)
Salary Rank (squared)	0.0002***
	(0.00002)
Constant	2.338***
	(0.143)
N	202
R^2	0.936
Adjusted R^2	0.936
Residual Std. Error	0.709 (df = 199)
F Statistic	1,461.332*** (df = 2; 199)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 10: Weighted ICT fit-line regression for lags

	Dependent variable:
	log(Weighted ICT)
Salary Rank	0.006***
	(0.001)
Constant	−0.246***
	(0.065)
N	202
R^2	0.353
Adjusted R^2	0.350
Residual Std. Error	0.476 (df = 200)
F Statistic	109.121*** (df = 1; 200)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Table 11: Weighted WhoScored fit-line regression for lags

	Dependent variable:
	log(Weighted WhoScored)
Salary Rank	0.009***
	(0.0002)
Constant	0.653***
	(0.019)
N	202
R^2	0.939
Adjusted R^2	0.939
Residual Std. Error	0.139 (df = 200)
F Statistic	3,074.695*** (df = 1; 200)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.