The Effect of Quotas on Wages: The Homegrown Player Rule in the Premier League*

Kate Ellison

Abstract

Quotas have long been used to increase representation of certain groups, especially in high-ranking positions. Theoretically, through classic models of labor supply and demand, they can also increase wages for workers helped by the quota. In this paper, I look at whether a foreign player restriction in the English Premier League impacts player wages and transfer fees. I find that this quota has no effect on the average of domestic players in the Premier League, although it does increase domestic player wages at clubs bound by the constraint of the quota. The quota also decreases domestic player transfer fees.

Introduction

The question of how quotas impact worker employment and compensation is of key importance from a policy standpoint. Quotas are often touted as a tool to improve workplace diversity and place minorities in high-ranking positions, but more research is needed on whether they merely improve representation, or if they also improve outcomes for workers included in the quota. There is a large body of literature exploring how employment and wages respond to policy changes such as minimum wages, parental leave, anti-discrimination laws, and many others (Card and Krueger, 1994; Antecol et al., 2018; Agan and Starr, 2018). Much of this work exploits a difference-in-differences design to compare the effect before and after a policy intervention. However, research exploring the wage effects of quotas specifically is much sparser. Existing literature has found some evidence for a modest increase in wages for women and workers with disabilities whose positions are protected or created due to quotas (Bertrand et al., 2019; Szerman, 2022). However, the Premier League quota provides a somewhat different context, as it was designed to protect the declining representation of a former majority group, rather than catapult underrepresented groups into higher positions. Therefore, the Homegrown Player Rule provides a compelling example of how quotas can impact wages in this somewhat more unusual case.

*I would like to thank Nina Roussille, Simon Jäger, Sara Ellison, Glenn Ellison, Lisa Ho, and Viola Corradini for their advice throughout this project. Additionally, I am very grateful to Jori Veng Pinje, not only for his generous provision of the wage data, but also for his helpful guidance. Finally, I would like to thank Neha Pant for her help in data scraping and Mackenzie Bivin for her help in developing my a theoretical model and for her moral support throughout.

In this paper, I ask the question of whether quotas increase worker wages. The example I explore is the 2010 adoption of the Homegrown Player Rule in the Premier League, England's top-tier professional soccer league, following the example of Kleven et al. (2013), which exploits data availability in professional soccer to answer more broadly applicable economic questions.

Implemented the start of the 2010-2011 season, the Homegrown Player Rule requires all Premier League clubs to register a minimum of eight homegrown players. If they cannot register eight homegrown players, they are penalized one spot on their roster per homegrown player they are short. Considering the twenty-five person roster, this amounts to a restriction of a maximum of seventeen non-homegrown players per roster. In order for a player to qualify as homegrown, they must have been trained at an English or Welsh soccer club for at least three years between the ages of fifteen and twenty-one. The Homegrown Player Rule places no restrictions based on nationality. Henceforth, however, I will be considering the impact of this rule on domestic as opposed to homegrown players for reasons outlined in my Empirical Methods section.

Player transactions in professional soccer occur in two different forms. First, clubs pay wages to their contracted players. When a player signs for a club, they are usually given a contract that lasts for a few years, paying a certain fixed wage across that period. Players will then re-negotiate new contracts with their club before their previous contract runs out, with new contracts often leading to higher wages. Second, two clubs carry out transactions in order for a player to move from one club to the other. This typically means the buying club will pay a negotiated fee to the selling club, though it is not uncommon for free transfers to occur, meaning no transfer fee is paid.

As my primary interest, I use a triple differences (DDD) empirical strategy to analyze the impact of the Homegrown Player Rule on domestic players' wages. Finding a null result on overall wages for domestic players, I then consider potentially differential for clubs for which the Homegrown Player Rule was binding versus nonbinding, again using a DDD design. I find that the Homegrown Player Rule increased wages for domestic players at clubs constrained by the Homegrown Player Rule quota. To supplement this analysis, I again use a DDD design to look at the impact of this rule on player transfer fees.

I will first build a simple theoretical framework to model the potential effects of a quota like the Homegrown Player Rule. Then, I describe my data and empirical methods, and discuss the necessary assumptions. Finally, I present the results of my analysis, draw conclusions, and propose potential areas for further research.

Theoretical Framework

In order to model the effects of a quota, I propose a simple framework to demonstrate how foreign and domestic soccer players' wages and transfer fees could be affected by a restriction on the number of foreign players. For simplicity, I assume that foreign and domestic players have the same labor supply, that is: $L_{s,dom} = L_{s,for} = L_s$. Then, I suppose that a team has demand $L_{d,dom}$ for domestic players and $L_{d,for}$ for foreign players. I also assume that, before the Homegrown Player Rule, the team is constrained by its roster limit, meaning it would prefer to employ more than twenty-five total foreign and domestic players, but it cannot

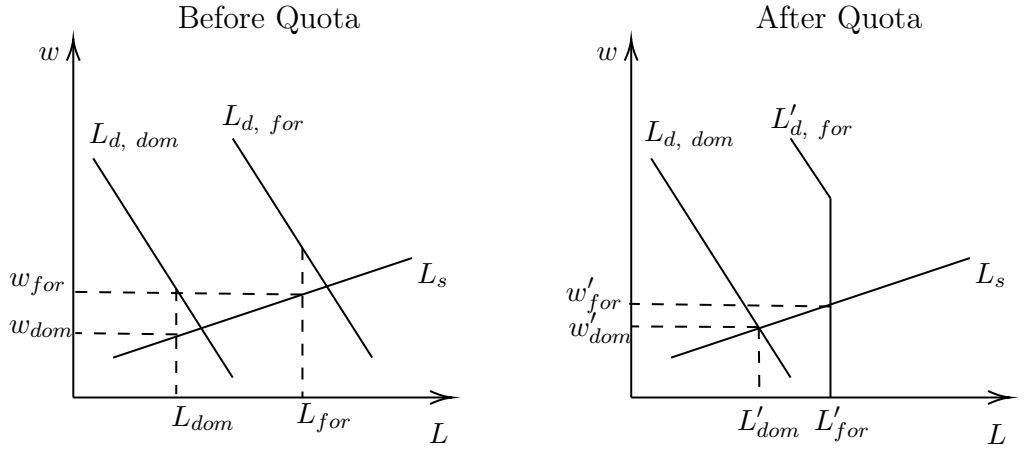

Figure 1: Demand for Foreign and Domestic Workers after a Quota

(though I later relax this assumption). Furthermore, the team will be constrained by the Homegrown Player Rule, meaning it currently employs more foreign players than the quota will allow. Then, its situation is modeled on the left of Figure 1.

In this case, the team chooses the optimal number of foreign and domestic players by equating the marginal benefit of hiring one more foreign or domestic player. However, the team is unable to hire up until the marginal benefit is zero, meaning they pay below the marginal product of labor for both foreign and domestic labor. Wages are instead determined by the labor supply curve, as indicated in Figure 1.

After the Homegrown Player Rule, $L_{d,for}$ becomes completely inelastic at the limit of seventeen foreign players. The team hires at $L'_{for} = 17 < L_{for}$. Because the team was previously constrained in the number of domestic players, it can and will want to now hire more, so $L'_{dom} > L_{dom}$. In Figure 1, the team is shown hiring up to the intersection of the domestic labor supply and demand curves, though this may not be the case (the team would not hire beyond that intersection). As we see in Figure 1, the model predicts that $\Delta w = w_{dom} - w_{for}$ will increase after the quota, meaning relative wages of domestic players will increase.

If we relax the assumption that the team's L_{dom} is constrained pre-quota, then even if w_{dom} is unchanged, Δw will increase due to the decrease in w_{for}. If the team's L_{for} is not constrained pre-quota, then we have the case where the quota is not binding. While this simple theoretical model predicts no effect on unconstrained teams because they are already hiring at their optimum, there may nevertheless be some effects. For example, their best domestic players might be tempted away by higher wages at previously constrained clubs, leading them to substitute towards the now-in-lower-demand foreign players.

To consider the impact of the quota on transfer fees, we assume that teams will only pay transfer fees for players whose marginal product of labor is greater than the wages they demand. Then, transfer fees are determined by some function of the player's marginal product to the two teams, along with the teams' bargaining power. Under these assumptions, average transfer fees for foreign players will weakly increase as a result of the quota because the average difference between foreign players' marginal products and wages demanded increases. On the other hand, transfer fees for domestic players will weakly decrease because the average

difference between their marginal products and wages will decrease.

Both of these model's predictions are borne out in the data, as shown below. Domestic players' transfer fees decrease across the board in the Premier League after the Homegrown Player Rule's implementation. Domestic wages do not significantly increase overall, but wages for domestic players do increase at clubs where the Homegrown Player Rule quota is binding.

Data

In order to analyze the empirical effect of the Homegrown Player Rule, I use three separate datasets. Though professional soccer players' wages are often discussed in the media, they are not required to be published by clubs, and thus it is difficult to find reliable wage data. In order to observe impact on player wages, I obtained a dataset taken from the video game Football Manager.[1] Football Manager wages are compiled from a variety of sources, but are not provided directly from clubs or players, so they should be seen as estimates. However, the accuracy of Football Manager data has proven to be relatively high when validated against external sources (Veng Pinje, 2013).

The wages dataset contains data from the 2005-2006 season to the 2011-2012 season in the top five European leagues (Premier League in England, La Liga in Spain, Ligue 1 in France, Serie A in Italy, and Bundesliga in Germany). These five leagues are widely considered to be the highest-level leagues in the world. Each observation in the wage dataset corresponds to a player-season, containing data for each player in one of the top five leagues in a particular season. The variables included are player name and demographics, club and league details, and wages. For each player, gross weekly wages are given in Pounds sterling, including any guaranteed bonuses in players' contracts. It should be noted, however, that conditional bonuses are also common in professional soccer contracts, although those are not included in the Football Manager dataset. I do not have a separate dataset on player bonuses, so this is a potential area for bias in the wage data.

The player demographics include nationality, and some players have multiple nationalities. For those players with multiple nationalities, I consider them to be domestic players in any country corresponding with one of their nationalities, although they would almost certainly not qualify as homegrown in all of those countries. Additionally, each player has a unique ID number within the dataset, allowing me to track players throughout the time period of 2005-2012. Summary statistics on wages in the Premier League and the other top four European leagues are given in Table 1.

In contrast to wages, transfer fees are usually published by clubs. I was able to download a dataset on transfer fees which had been scraped from Transfermarkt.us, a website that tracks many different variables related to soccer players, including their transfer fees. I was unable to find specifics on how Transfermarkt determines and verifies transfer fees, although it is widely regarded as a trusted source for soccer data.

This dataset contains data on all transfers to and from teams in any of the top five European leagues from 1992-2022, though I restricted my analysis to 2005-2015 in order to reduce the impact of other confounding events, such as Brexit. The transfers dataset

[1] I am very grateful to Jori Veng Pinje for providing me with these data.

contains one observation per transfer event, meaning a transfer of a player between two clubs. For each transfer, the dataset records the player's name, both clubs' names, the direction of the transfer (in vs out), the type of transfer (loan vs permanent move), some player characteristics, and the transfer fee paid by the buying club. For permanent transfers, the transfer fee is typically published by at least one of the involved clubs, though if both clubs agree, the fee may be undisclosed. Reasons for not disclosing a transfer fee could include concern about fan reaction, jealousy from other players, and specific club preferences, all of which could potentially bias the reported transfer fees. Overall, around 15% of permanent transfers' fees are unreported.

However, the transfer dataset does not include data on player nationalities. In order to compare transfer fees of foreign versus domestic players, I needed player nationalities. Because the wages dataset only covers the period from 2005-2012, not all players in the transfer dataset are covered by the wage dataset, so I could not use nationality data from the wage dataset. I was able to scrape nationality data from FBRef.com, which lists the nationality of all top professional players, male and female, from the 1950s onwards. This is a simple cross-sectional dataset with one observation per player, and each observation listing just the player's name and nationality. According to FBRef, this dataset was compiled from players' FIFA-registered nationalities. FIFA is the main governing body for global football. As a result, each player is listed with only one nationality, which is thought of as their "main" nationality.

To link these datasets, I match player names from the nationality dataset with observations in the transfer dataset. Many players show up multiple times in the transfer dataset, representing multiple transfers so the player's nationality was applied to all observations. One complication in merging the datasets is the non-uniqueness of names in the nationality dataset. For each duplicate name, if all duplicates have the same nationality, I keep the name in the dataset. If duplicate names have different nationalities, then I drop them, because I will not be able to later make the distinction of whether that player is domestic or not in their league.

In total, around 3% of players in the nationality dataset were removed because of other players of their same name from different countries. This figure seems unusually high, but it is largely driven by many soccer players' preferences to go by a single name as opposed to a first and last name. It should also be noted that the majority of the players in the nationality dataset never played in the top five European leagues and would not be included in final analysis regardless. When merging nationality and transfer data, of the roughly 100,000 transfers in the transfer dataset, 18% of them were unable to be matched with a player in the nationality dataset, largely driven by small differences in spelling, accents, punctuation, and nicknames. There could be cause for some concern here if names that are unmatched are correlated with certain nationalities.

Finally, in order to remove duplicate transfer records (from a player moving from one club in the top 5 leagues to another), I kept only transfers in the "in" direction (roughly 50% of transfers). Table 2 presents summary statistics for the transfers that are able to be used in regression (i.e. they have known transfer fees).

Empirical Methods

Triple Differences

I use a triple differences identification strategy to estimate the impact of the Homegrown Player Rule on wages and transfer fees. This consists of comparing an outcome variable (player wage or transfer fee) between domestic and foreign players, before and after the Homegrown Player Rule, between treated and untreated groups. Without considering for which clubs the Homegrown Player Rule is binding, the Premier League is the only treated league, while the other four top European leagues serve as control leagues. I selected these four European leagues as controls mostly due to similarity in quality and popularity, but it is important to note as well that the Premier League is unique in its global popularity and financial wherewithal. Therefore, these four leagues cannot be seen as perfect controls for the Premier League in absence of a Homegrown Player Rule.

In my regressions comparing Premier League clubs that are constrained and unconstrained by the Homegrown Player Rule, the treated group represents players at clubs at which the Homegrown Player Rule was binding in the 2008-2009 season (the last control season). The control group represents players at clubs that were not constrained by the Homegrown Player rule in 2008-2009.

Due to my triple differences strategy, the regression has many interaction terms, with β_7 being the coefficient of interest.

$$\begin{aligned} y_{it} = \beta_0 &+ \beta_1(\text{Treated}_{it}) + \beta_2(\text{Post}_t) + \beta_3(\text{Domestic}_{it}) \\ &+ \beta_4(\text{Treated}_{it} \times \text{Post}_t) \\ &+ \beta_5(\text{Treated}_{it} \times \text{Domestic}_{it}) \\ &+ \beta_6(\text{Post}_t \times \text{Domestic}_{it}) \\ &+ \beta_7(\text{Treated}_{it} \times \text{Post}_t \times \text{Domestic}_{it}) + \beta X_{it} + \gamma_t + \delta_i + \epsilon_{it} \end{aligned}$$

In this regression, y_{it} represents an outcome of either player wage, transfer fee, log wage, or log transfer fee, depending on the regression specification. Treated_{it} represents whether an individual i is in the treatment group (as described above) at time t, Post_t represents whether the observation takes place before or after 2009, and Domestic_{it} represents whether or not player i is counted as a domestic player in their league at time t. X_{it} represents a vector of control variables, including player age, age^2, and league and club fixed effects. γ_t represents time fixed effects, and δ_i represents player fixed effects.

Note that, in this regression, I am comparing *domestic* and *foreign* as opposed to *homegrown* and *non-homegrown* players. This is in order to better compare between treatment and control leagues, because in leagues other than the Premier League, there does not exist a homegrown designation. I found it too arduous a data collection task to determine which players in each leagues had been trained in that country for at least three years before they turned twenty-one. This is certainly an area that warrants further research, as the effect on homegrown players is plausibly different from the effect on domestic players, who may or may not qualify as homegrown. When classifying players as domestic, I count players of both Welsh and English nationality as UK domestic players.

Additionally, I classify observations from the 2009-2010 season onward as post-treatment due to concern over anticipatory effects of the Homegrown Player Rule. These effects can be

seen in Figure 3 where there is a significant dip in foreign player transfer fees in the Premier League in 2009. The Homegrown Player Rule was announced in 2009, but did not go into effect until a year later at the start of the 2010-2011 season. Therefore, it is reasonable that clubs would have made transfer and contract decisions in 2009 with the future rule in mind. The result of this regression can be interpreted as the bonus (either in wages or transfer fees) given to domestic players over foreign players as a result of this exogenous quota. Of course, this empirical strategy is only valid if the assumptions underpinning a triple differences design are satisfied.

For the constrained treatment group, clubs are included if they had fewer than eight domestic players in the 2008-09 season. The histogram of domestic players per club in 2008-09 is plotted in Figure 7. There are four clubs (out of twenty) in the treatment group: Arsenal, Chelsea, Liverpool, and Sunderland.

I would have also liked to include a quadruple differences regression comparing domestic vs foreign players at constrained vs unconstrained clubs in the Premier League vs other four leagues before vs after the Homegrown Player Rule. However, the key issue with this regression is the definition of constrained vs unconstrained in leagues other than the Premier League. In order to be constrained in the Premier League, a club must have fewer than eight domestic players in the 2009-10 season. When this same definition is applied to the other four leagues, there is only one club (out of 98) that fits this definition, as the Premier League is the most cosmopolitan league. Using this one club as a control group would not be very meaningful, meaning the results of this quadruple differences regression would not be intuitively interpretable.

Parallel Trends

In a triple differences empirical method, we require one parallel trends assumption, which can be written as:

$$\begin{aligned}
&(E[y_{0,it}|\text{Treated}_{it}=1,\text{Domestic}_{it}=1,\text{Post}_t=1]-E[y_{0,it}|\text{Treated}_{it}=1,\text{Domestic}_{it}=1,\text{Post}_t=0])\\
&-(E[y_{0,it}|\text{Treated}_{it}=1,\text{Domestic}_{it}=0,\text{Post}_t=1]-E[y_{0,it}|\text{Treated}_{it}=1,\text{Domestic}_{it}=0,\text{Post}_t=0])\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad =\\
&(E[y_{0,it}|\text{Treated}_{it}=0,\text{Domestic}_{it}=1,\text{Post}_t=1]-E[y_{0,it}|\text{Treated}_{it}=0,\text{Domestic}_{it}=1,\text{Post}_t=0])\\
&-(E[y_{0,it}|\text{Treated}_{it}=0,\text{Domestic}_{it}=0,\text{Post}_t=1]-E[y_{0,it}|\text{Treated}_{it}=0,\text{Domestic}_{it}=0,\text{Post}_t=0])
\end{aligned}$$

as shown in Olden and Møen (2022). In this equation, $y_{0,it}$ represents the outcome of wages or transfer fees if treatment had never occurred. In words, this refers to a counterfactual world in which the Homegrown Player Rule was not implemented. Then, the difference-in-differences of foreign versus domestic and pre- versus post-2009 outcomes in the treatment group should equal the the difference in differences of foreign versus domestic and pre- versus post-2009 outcomes in the control group. In order to assess the validity of these assumptions, I look at the pre-trends for foreign versus domestic players in the Premier League and control leagues.

In Figure 2, the average wages of foreign and domestic players are plotted in the Premier League and the four control leagues. We see that the pre-trends seem roughly parallel up through the 2008-2009 season, the season before the Homegrown Player Rule was announced.

The only year that suggest some cause for concern is 2006, when the trends in the treatment group are somewhat divergent to trends in the control group.

Similarly, I plot the same averages for transfer fees in Figure 3. The parallel trends assumption is somewhat less compelling for transfer fees. The gap between foreign and domestic players' wages appears to widen before 2009 in the control leagues, while this does not occur in the Premier League. Interestingly, in the Premier League, domestic players' fees are higher than foreign players' fees before 2010, but lower after. In contrast, in the four control leagues, foreign players' transfer fees are consistently higher than those of domestic players.

Another cause for concern in these graphs is the significant but temporary drop in foreign players' transfer fees in the control leagues from 2010-2012. Overall, the volatility of transfer fees across the period from 2005-2015 could suggest that there are other uncontrolled factors impacting transfer fees, like financial regulations, taxation, or selection into publication of transfer fees.

Spillover Effects

Another key assumption to a triple differences empirical method is that the treatment and control groups are independent, and that any treatment applied to the treatment group has no effect on the control group. In the case of the Homegrown Player Rule, we might be concerned that the rule would have spillover effects into the four control leagues. For example, if clubs in the control leagues employed many English players, and their wages increased as a result of better outside options and more demand for English players, this would increase wages on average for foreign players in the control leagues, biasing the DDD estimate.

However, in this case, I do not believe that spillover effects should cause great concern. There is very little mobility of English players to the control leagues. From the 2005-06 season to the 2011-12 season, there are only 33 player-season observations of English or Welsh domestic players employed by control league clubs, compared to a total of 1458 English or Welsh domestic player-season observations total. Additionally, using the transfer data, there were only eight transfers of English or Welsh domestic players to clubs in the control leagues from 2005-2015. Due to these small numbers, I do not believe that spillover effects will significantly disrupt the independence of the treatment and control groups. Full statistics for the origin clubs of UK domestic players transferring to Premier League clubs are given in Table 3.

Nevertheless, I also provide results from a simple differences-in-differences regression comparing UK domestic players and all other players before and after the Homegrown Player Rule in Table 8. Of course, this regression loses the dimension of the Premier League vs control leagues, so it does not account for the general trend of globalization and increased hiring of foreign players across leagues. However, its results may be less threatened by potential spillover effects. The regression equation is given by:

$$y_{it} = \alpha_0 + \alpha_1(\text{Post}_t) + \alpha_2(\text{Domestic}_{it}) + \alpha_3(\text{Post}_t \times \text{Domestic}_{it}) + \beta X_{it} + \gamma_t + \delta_i + \epsilon_{it}$$

where α_3 is the coefficient of interest, and all variables are defined as in the DDD regression.

Results

First, I look at evidence for whether the Homegrown Player Rule increased representation of domestic players in the Premier League. Results are presented in Table 4. I find that, in absolute terms, the number of domestic players at Premier League clubs does not increase after the Homegrown Player Rule. This should not be too surprising, though, considering that the Homegrown Player Rule was intended to slow the trend of increasing numbers of foreign players in the Premier League, not completely reverse it. When compared to the four control leagues, there still is no relative increase in the number of domestic players in the Premier League. However, when I compare constrained and unconstrained clubs within the Premier League, I find that the quota does indeed increase the number of domestic players at constrained clubs.

In Table 5, I present the results for the triple differences regression of the effect of the Homegrown Player Rule on wages for all players in the Premier League. In regressions (1) - (4), log weekly wages are regressed on the full set of interaction terms in the regression given above, as well as age controls and certain fixed effects, indicated for each regression. I find that there is no effect of the Homegrown Player Rule on domestic players' wages, given that the magnitudes of the coefficients of interest are small (between -4% and -8%) and none of them are statistically significant at the 5% level. In Figure 4, I plot the DDD coefficients by year on log wages for domestic Premier League players, based on Regression (4) from Table 5 including the full set of controls. There appears to be a slight downward trend in wages after the Homegrown Player Rule, but the only season with significantly lower wages than the 2008-2009 season is the 2011-2012 season.

In Table 6, I present the results from regressions measuring the effect of the Homegrown Player rule on wages of domestic players at constrained teams. For Regression (1), I run a simple difference-in-differences regression comparing domestic player wages at constrained vs unconstrained clubs before and after the Homegrown Player rule. I estimate that wages of domestic players at constrained clubs increases by 36% when compared to wages of domestic players at constrained clubs. Regressions (2)-(4) are triple differences regressions comparing wages of domestic vs foreign players at constrained vs unconstrained clubs before and after the homegrown player rule. In Regression (4), with all controls included, I find that domestic player wages at constrained clubs rose by 42% after the Homegrown Player rule. In Figure 5, I plot the coefficients on log wages of domestic players at constrained clubs using the same controls as in Regression (4) from Table 6. We see that wages for these domestic players increase significantly in the first treatment year and remain well above pre-treatment levels. It should be noted that while these effects are large and significant, they are also imprecisely measured, with standard errors around 10-15 percentage points.

To supplement my analysis of player wages, I use the same triple differences strategy to analyze the impact on transfer fees. Table 7 presents the results of these regressions. As with wages, I regress transfer fees on the full set of interaction terms from my regression equations. Regressions (1) - (4) are performed using transfer fees. However, Regression (5) provides the log regression of 1 + transfer fee in order to still include the modal transfer fee of 0. All regressions (1)-(4) show a large decrease in transfer fees of domestic players, between €1.6 million and €1.9 million. Regression (4) excludes the 99th percentile of transfer fees. Transfer fees have a right-skewed distribution, meaning that a few big-ticket transfers before or after

the 2010 Homegrown Player Rule implementation could potentially bias estimates. However, we see in regression (4) that this does not appear to be the case, as the DDD coefficient does not change much. The main difference is that the standard error is lower because the standard deviation of transfer fees is lower. All of these estimates are significant at the 5% level, and the coefficient in regression (4) is significant at the 1% level. In Regression (5), which runs the DDD regression on log transfer fees, the effect is also very large, with the Homegrown Player Rule estimated to reduce transfer fees by 30%. In Figure 6, just as with wages, I plot the coefficients from the DDD regression of transfer fees. We see a steady decline in transfer fees after the implementation of the rule.

Conclusion

While I do not find an overall positive wage effect on domestic players, I do find that the Homegrown Player Rule increased wages for domestic players at clubs constrained by the rule. This result offers some promising evidence for the ability of quotas to not only increase representation, but also increase wages of workers protected by the quota. The zero effect on overall wages for domestic players in the Premier League, though, may be cause for concern. One potential explanation could be that the best domestic players simply moved from unconstrained to constrained clubs who were looking to replace their talented foreign players. Additionally, wages are sticky, especially in soccer where wages are given in multi-year contracts. My analysis of wages only considers wages up through 2012, which may not allow time for wages to fully adjust after the implementation of the rule.

The decrease in transfer fees, while also consistent with model predictions, should not be accepted unequivocally either. Selection into transfer fee publication is a potential concern, as well as the volatility of transfer fees across seasons.

In the case of both wages and transfer fees, it is important to recognize that neither are complete measures of player compensation, nor do they encompass all transactions that occur in the professional soccer market. Just to name a few: professional soccer players have sponsorship deals with companies; player and club agents receive sizable commissions from transfer fees, and those figures are rarely publicized; finally, differential tax rates for foreign and domestic players mean that net wages for foreign and domestic players may be significantly different from their gross wages, which are used in this analysis.

In the case of the Homegrown Player Rule, questions still remain. Potential areas for future research include looking at more recent data to estimate longer-term effects on wages, as well as developing an empirical strategy to distinguish the effect on wages and transfer fees of all homegrown players, not just domestic players.

More broadly, the effect of quotas on wages is by no means empirically certain. While representation is an important metric, at the end of the way, workers really care about wages. In addition to increasing representation, one of the express goals of quotas should be to improve compensation for workers covered by the quota. In order to assess the impact of quotas on workers, it is necessary to examine what happens to relative wages as a result of quotas.

References

Agan, Amanda and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *Quarterly Journal of Economics*, 2018, *133* (191), 235.

Antecol, Heather, Kelly Bedard, and Jenna Stearns, "Equal but Inequitable: Who Benefits from Gender-Neutral Tenure Clock Stopping Policies?," *American Economic Review*, 2018, *108* (9), 2420–2441.

Bertrand, Marianne, Sandra E Black, Sissel Jensen, and Adriana Lleras-Muney, "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway," *The Review of Economic Studies*, 2019, *86* (1), 191–239.

Card, David and Alan Krueger, "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," *American Economic Review*, 1994, *84* (4), 772–793.

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

Olden, Andreas and Jarle Møen, "The triple difference estimator," *The Econometrics Journal*, 2022, *25*, 531–553.

Pinje, Jori Veng, "Wage- and Tax-Induced Mobility of Footballers among 30 European Countries," 2013.

Szerman, Christiane, "The Labor Market Effects of Disability Hiring Quotas," 2022.

Figures and Tables

Table 1: Wages in the Top 5 European Leagues, 2005-2012

	Premier League		Four Control Leagues	
	Pre-2009	Post-2009	Pre-2009	Post-2009
Number on payroll	1833	1390	7153	5852
Mean	20.6	30.8	12.6	17.9
Standard Deviation	20.8	28.4	13.7	24.2
Minimum	0	0	0	0
Maximum	160	180	169	283

This includes all estimated weekly wages paid to players in the top 5 European leagues from 2005-2012. Mean, standard deviation, minimum, and maximum are given in thousands of Pounds.

Log weekly wages (in GBP) of domestic and foreign players are plotted from 2006-2012. Each year on the x-axis corresponds to wages during the season ending in that year (i.e. 2009 refers to the 2008-2009 season).

Figure 2: Wages Trends Premier League and Control Leagues

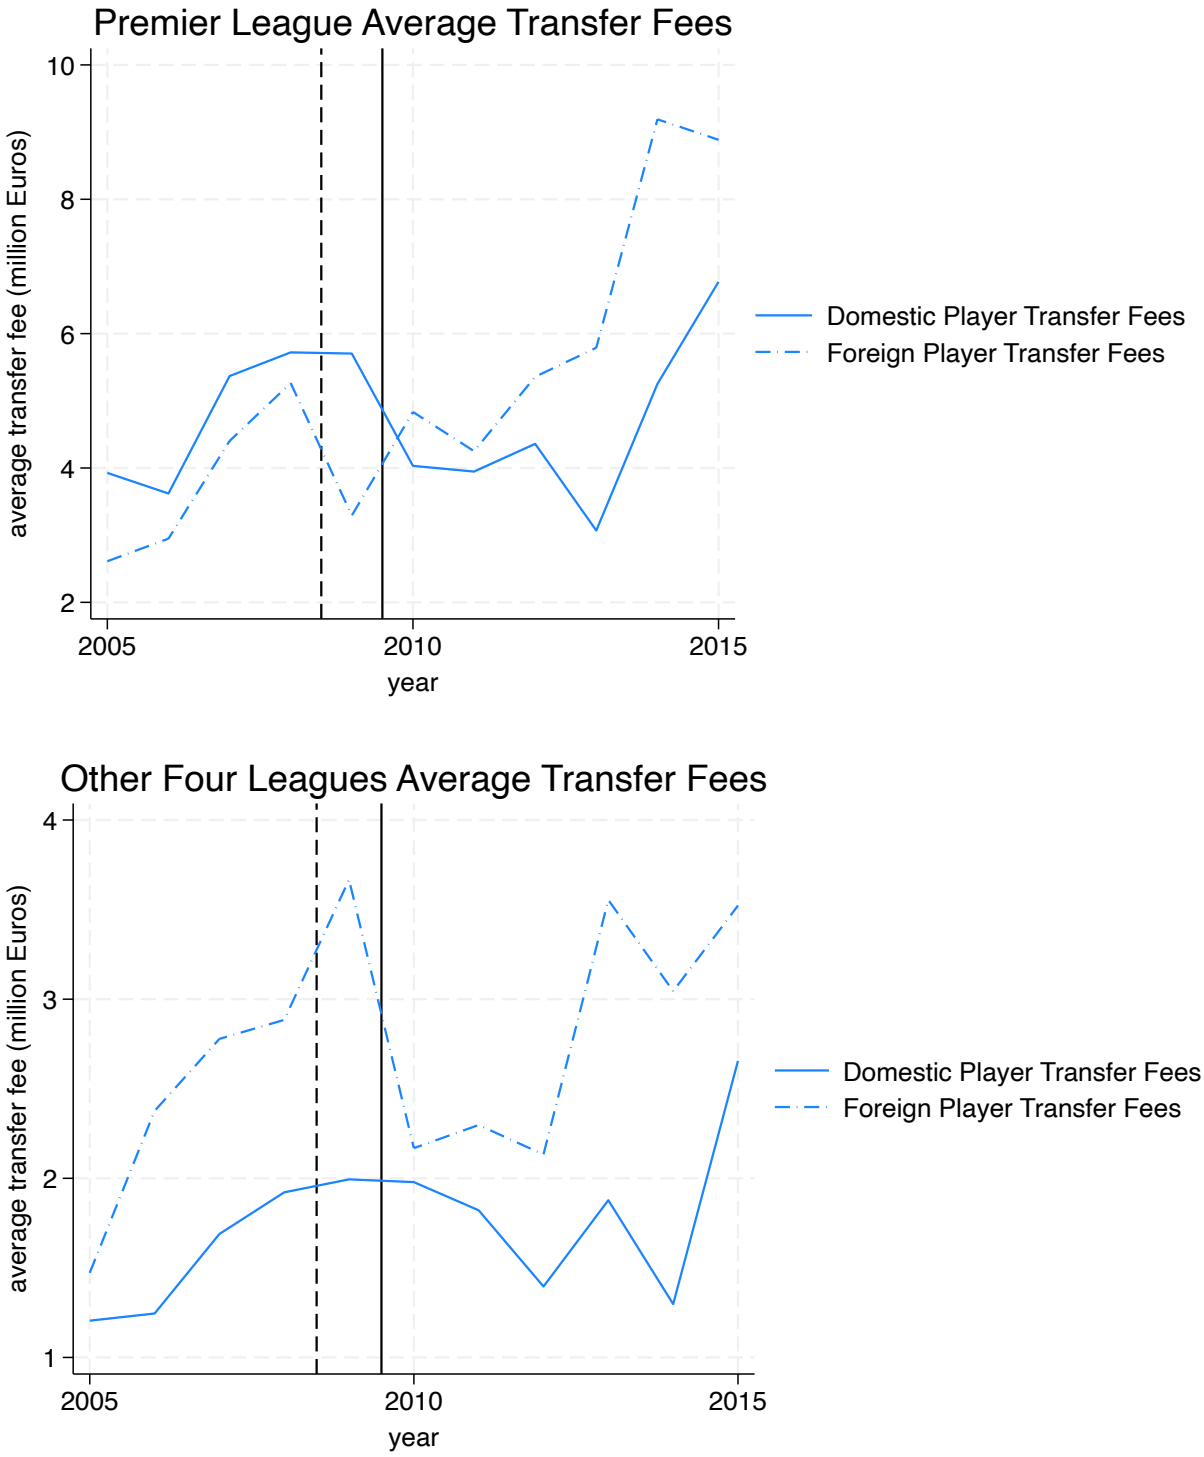

Average transfer fees of (in millions of Euros) of domestic and foreign players are plotted from 2005-2010. Each year on the x-axis corresponds to transfers occurring during that calendar year.

Figure 3: Transfer Fee Trends Premier League and Control Leagues

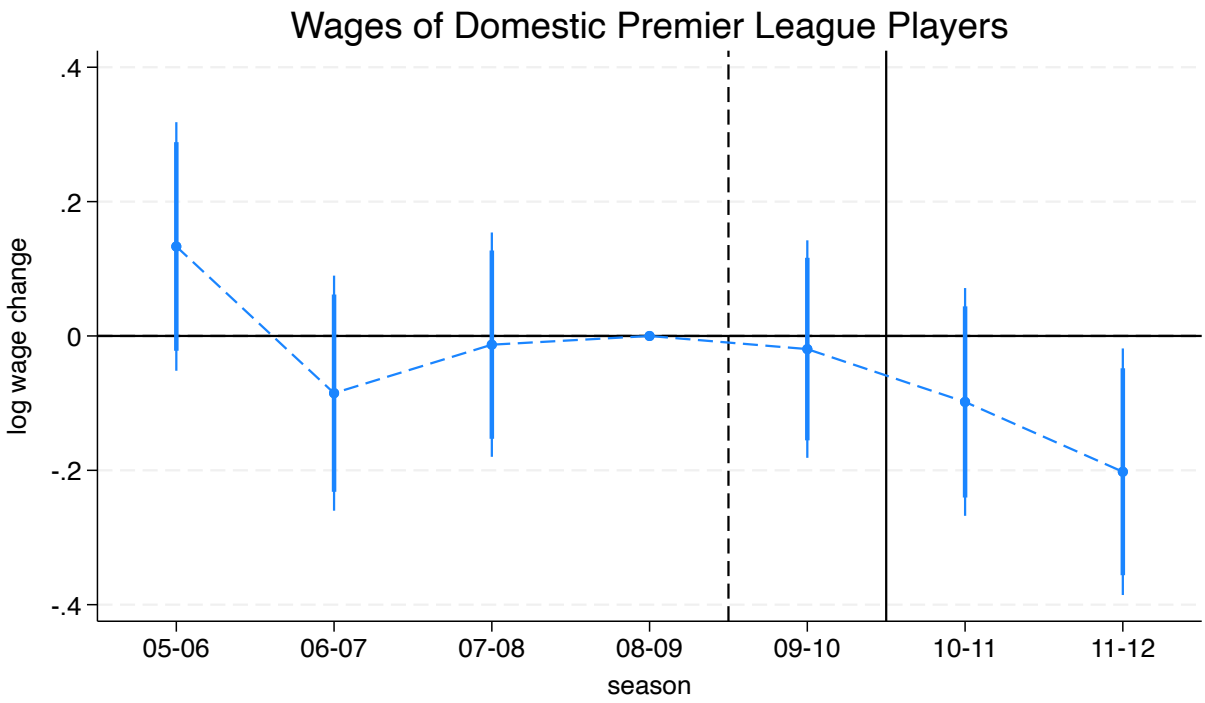

The coefficients on log wages for domestic Premier League players are plotted, compared to foreign Premier League players and control league players. These coefficients are derived from a regression that includes age controls, and league, club, and player fixed effects. The bars represent 90% and 95% confidence intervals.

Figure 4: Change in Premier League Domestic Player Wages

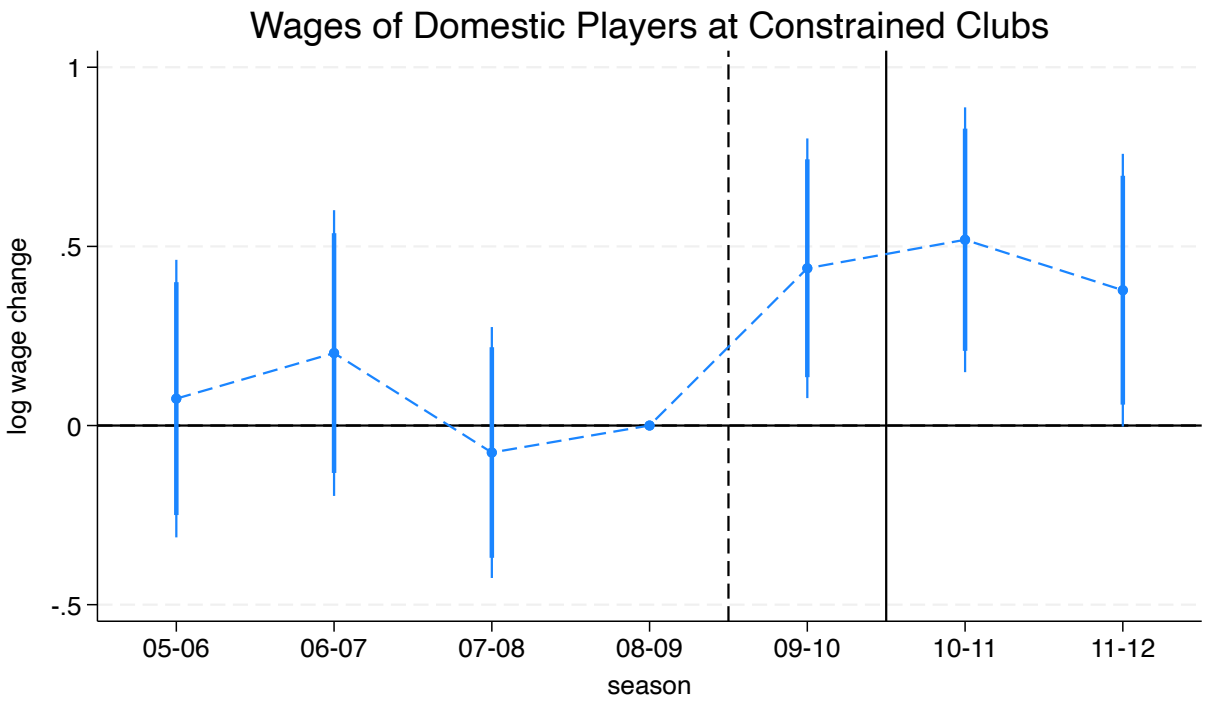

The coefficients on log wages for domestic players at constrained clubs are plotted, compared to foreign players and players at unconstrained clubs. These coefficients are derived from a regression that includes age controls, and league, club, and player fixed effects. The bars represent 90% and 95% confidence intervals.

Figure 5: Change in Premier League Constrained Domestic Player Wages

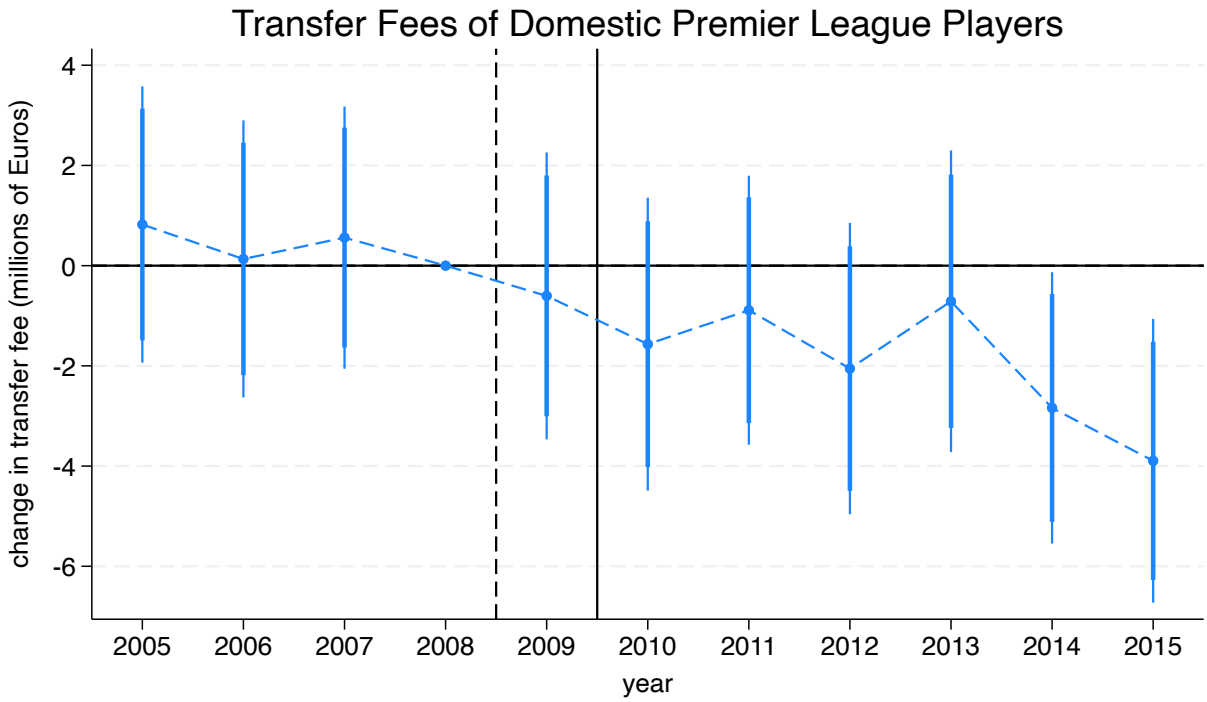

The coefficient on transfer fees for domestic Premier League players are plotted, compared to foreign Premier League players and control league players. These coefficients are derived from a regression that includes age controls, and league and club, and player fixed effects. The bars represent 90% and 95% confidence intervals.

Figure 6: Change in Premier League Domestic Transfer Fees

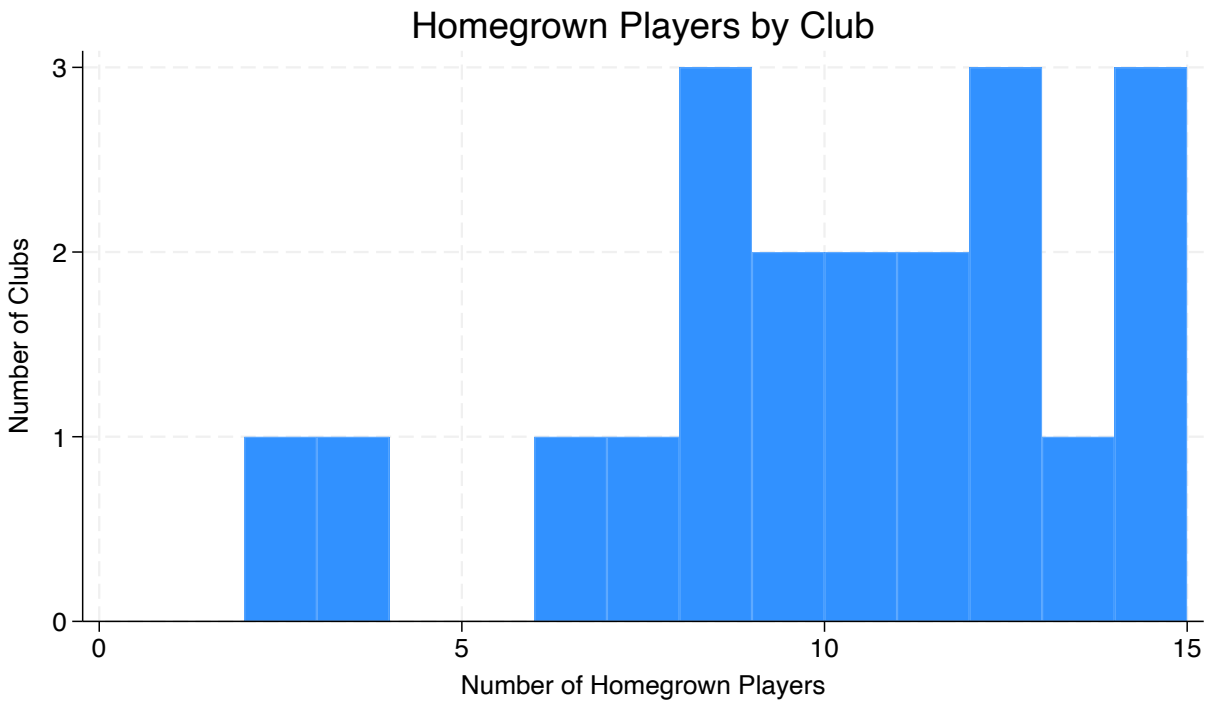

This histogram plots the number of UK domestic players at each Premier League club in the 2008-2009 season (the last control season). While domestic is not a perfect analog for homegrown only one player in the league that year would have qualified as homegrown but not domestic, so this is a pretty accurate representation of which clubs were constrained by the Homegrown Player rule.

Figure 7: Number of Domestic Players At Premier League Clubs in 2008-09

Table 2: Transfer Fees in the Top 5 European Leagues, 2005-2015

	Premier League		Four Control Leagues	
	Pre-2009	Post-2009	Pre-2009	Post-2009
Number of transfers	603	952	1947	3715
Mean # of transfers/year	151	136	487	531
Mean	4.1	5.7	2.0	2.4
Standard Deviation	6.1	8.9	3.6	5.8
Minimum	0	0	0	0
Maximum	43.9	76	35.5	101

This includes all transfers into the Premier League and the Control Leagues from 2005-2015 for which the transfer fee was known. Mean, standard deviation, minimum, and maximum are given in millions of Euros.

Table 3: Origins of UK Domestic Players Transferring to Premier League Clubs, 2005-2015

Club transferring from	Number	Percent
Other English club	417	95
English youth team	8	2
Club in other 4 European leagues	5	1
Club in other foreign league	8	2
Total	438	100

This includes all transfers into the Premier League and the Control Leagues from 2005-2015 for which the transfer fee was known. Other English club refers to any club in the English and Welsh football pyramid, including teams in the Premier League and lower league. English youth team refers to the youth team of any English club. A club in the other 4 European leagues refers to clubs in La Liga, Ligue 1, Bundesliga, or Serie A.

Table 4: Number of Domestic Players at Premier League Clubs

	(1)	(2)	(3)
Number of domestic players	-26.91*	-42.67	32.75**
	(14.85)	(38.71)	(14.83)
Regression	DID	DDD	DDD
Constrained?			Yes

Regression (1) is the DID regression of number of domestic vs foreign players before and after the HG player rule. Regression (2) is the DDD regression of number of domestic vs foreign players before and after the HG player rule in the Premier League vs the control leagues. Regression (3) is the DDD regression of the number of domestic vs foreign players before vs afteer the HG player rule at constrained vs unconstrained clubs

*** = significant at the 1% level, ** = significant at the 5% level, * = significant at the 10% level

Table 5: Triple Differences Regression of Wages

	(1)	(2)	(3)	(4)
$\text{Treated}_{it} \times \text{Post}_t \times \text{Domestic}_{it}$	-0.0517	-0.0703	-0.0665	-0.0643
	(0.0804)	(0.0562)	(0.0561)	(0.0547)
age	1.0626***	1.1206***	1.1151***	1.1067***
	(0.0241)	(0.0555)	(0.0553)	(0.0543)
age^2	-0.0180***	-0.0201***	-0.0201***	-0.0197***
	(0.0004)	(0.0006)	(0.0006)	(0.0006)
Year fixed effects?	Yes	Yes	Yes	Yes
Player fixed effects?		Yes	Yes	Yes
League fixed effects?			Yes	Yes
Club fixed effects?				Yes
R^2	0.34	0.82	0.82	0.83
Number of observations	16210	13714	13714	13714

In regressions (1)-(4) log weekly wages (in thousands of GBP) are regressed on all the terms from the above regression.

*** = significant at the 1% level, ** = significant at the 5% level, * = significant at the 10% level

Table 6: Constrained vs Unconstrained Regression of Wages

	(1)	(2)	(3)	(4)
$\text{Treated}_{it} \times \text{Post}_t \times \text{Domestic}_{it}$	0.3631***	0.8706***	0.4406***	0.4218***
	(0.1197)	(0.1986)	(0.1319)	(0.1342)
age	0.8797***	1.0618***	0.8863***	0.9042***
	(0.1604)	(0.0498)	(0.0973)	(0.0996)
age^2	-0.0184***	-0.0181***	-0.0172***	-0.0176***
	(0.0016)	(0.0009)	(0.0012)	(0.0012)
Regression	DID	DDD	DDD	DDD
Year fixed effects?	Yes	Yes	Yes	Yes
Player fixed effects?	Yes		Yes	Yes
Club fixed effects?	Yes			Yes
R^2	0.87	0.40	0.86	0.87
Number of observations	1192	3222	2658	2658

In regressions (1)-(4) log weekly wages (in thousands of GBP) are regressed on all the terms from the above regression.
*** = significant at the 1% level, ** = significant at the 5% level, * = significant at the 10% level

Table 7: Triple Differences Regression of Transfer Fees

	(1)	(2)	(3)	(4)	(5)
Transfer fee	Yes	Yes	Yes	Yes	
Log(1+TF)					Yes
$\text{Treated}_{it} \times \text{Post}_t \times \text{Domestic}_{it}$	-1.8603**	-1.8611**	-1.6307**	-1.7671***	-0.3140***
	(0.7903)	(0.7904)	(0.7143)	(0.5591)	(0.1044)
age	1.5884***	1.5586***	2.0003***	1.6259***	0.4095***
	(0.1354)	(0.1343)	(0.1437)	(0.1064)	(0.0219)
age^2	-0.0334***	-0.0330***	-0.0404***	-0.0331***	-0.0085***
	(0.0025)	(0.0025)	(0.0027)	(0.0020)	(0.0004)
Year fixed effects?	Yes	Yes	Yes	Yes	Yes
League fixed effects?		Yes	Yes	Yes	Yes
Club fixed effects?			Yes	Yes	Yes
Excluding 99th percentile?				Yes	
R^2	0.08	0.09	0.34	0.33	0.39
Number of observations	7217	7217	7217	7149	7216

In regressions (1)-(4) transfer fees or log transfer fees (in millions of Euros) are regressed on all the terms from the above regression.
*** = significant at the 1% level, ** = significant at the 5% level, * = significant at the 10% level

Table 8: Difference-in-Difference Regressions of Wages and Transfer Fees

	(1)	(2)	(3)
Log wage	Yes	Yes	
Transfer fee			Yes
$\text{Post}_t \times \text{Domestic}_{it}$	0.1084*	-0.0882**	-0.1957
	(0.0565)	(0.0370)	(0.7046)
age	1.1098***	1.1359***	1.6063***
	(0.0246)	(0.0554)	(0.1402)
age^2	-0.0189***	-0.0203***	-0.0338***
	(0.0005)	(0.0006)	(0.0026)
Year fixed effects?	Yes	Yes	Yes
Player fixed effects?		Yes	
R^2	0.29	0.81	0.04
Number of observations	16210	13714	7217

In regressions (1)-(4) log weekly wages (in thousands of GBP) or transfer fees (in millions of Euros) are regressed on the terms from the difference-in-differences regression.
*** = significant at the 1% level, ** = significant at the 5% level, * = significant at the 10% level