

# The Dynamics of Racial Discrimination in a Virtual Labour Market

Alex Bryson<sup>♠</sup>

University College London

Arnaud Chevalier<sup>♣</sup>

Royal Holloway, University of London

February 2023

(Preliminary and incomplete – Do not Quote)

## ABSTRACT

We test Becker's prediction on employer racial prejudice in a dynamic environment where discrimination is not illegal and information on all workers' productivity is available to all firms. The firms considered are the more than 2 million participants to Fantasy Football (soccer), an on-line game where participants manage a team of real footballers over a nine-month period. We find that the marginal non-white footballer needs to be more productive than a white worker to be recruited, promoted and to avoid dismissal. These decisions lead to firms becoming “whiter” over time.

These effects are strongest among the most productive footballers but, because elite footballers are close substitutes, the productivity losses of discrimination are limited. Allied to an environment where discrimination is not illegal, and mostly private, this allows virtual employers to satisfy their taste for discrimination. The results are robust to three different measures of race.

Key words: race; labour market discrimination; taste-based discrimination; football

JEL: J15; J23; J24; J71; M51

### *Acknowledgements*

The authors thank the Centre for Economic Performance for seed money for this research. We also thank Wiktor Piotrowski and Martin Guenther for excellent research assistance.

---

<sup>♠</sup> Bryson is also affiliated to NIESR (London), IZA (Bonn) and Rutgers (New Jersey): a.bryson@ucl.ac.uk

<sup>♣</sup> Chevalier is also affiliated to IZA (Bonn), Vive (Copenhagen), and ROA (Maastricht): arnaud.chevalier@rhul.ac.uk

## 1. INTRODUCTION

Despite legislation prohibiting differential treatment of workers by virtue of race, empirical evidence from around the world indicates that racial differences in hiring rates are large (Lippens et al., 2023) and persistent (Di Stasio and Heath, 2019). Since racial discrimination is unlawful, the evidence comes predominantly from correspondence studies whereby similar CVs differing only in (implied) ethnicity are sent as a response to a job advert (see Baert, 2018 for a review). While informative about employers' discriminatory behavior they suffer from at least two drawbacks. First, they focus on a single point in the hiring process so are unable to establish how employer discriminatory behaviour evolves over a period with new hires and fires. Second, they are unable to isolate the potential sources of discrimination.

We contribute to the literature on racial discrimination by examining hiring, firing and promotion decisions in a dynamic virtual labour market setting where we observe decisions in 38 periods spanning 9 months. For each period we observe hiring, firing and promotion decisions for all employers, and the universe of all potential employees.

Furthermore, the institutional features of our setting described below reduce the source of discrimination to employer taste-based discrimination, allowing us to identify its extent and dynamics, and the costs to the employer of indulging in their taste for discrimination.

We offer a novel strategy to explore discriminatory behavior in hiring, firing and promotion decisions in a virtual labour market - the Fantasy Premier League (FPL). FPL is an on-line management competition based on the Premier League, the foremost league

in English football (soccer), entered by more than 2 million individuals worldwide for the year of interest.<sup>1</sup> Participants to FPL behave like employers hiring their initial squad of professional footballers, subject to a fixed budget, and subsequently adjusting it on a weekly basis over nine months to maximise production. Production is only driven by managerial talent – i.e. selecting the most productive real footballers – and not by video game competencies.

In contrast to earlier studies, our data allow us to isolate employer taste-based discrimination. Participants to FPL have perfect information about the race and week-by-week on-field performance of every professional football player in the league so there are no grounds for statistical discrimination. In the absence of customers and co-worker relations we can rule these out as potential sources of discriminatory behaviour. Moreover, there is no self-selection of employees to firms, no strategic motivations when hiring and the prices of players (their transfer costs) are set exogenously outside the firm so are not subject to the racial preferences of managers. The selected team thus reflects the pure preferences of the participant. Moreover, in this environment discrimination is not unlawful (information regarding the workers employed remains private to the firm) and can be fully satisfied.

The setting approximates a laboratory experiment, in the sense that employers have identical information, objectives and constraints. This provides us with an observational study where we can observe weekly decisions over a nine-month period, including not only hiring, but also firing and (a proxy for) promotion. As well as hiring and firing footballers so as to optimise their team, participants have to select 11 players from a squad of 15 whose

---

<sup>1</sup> <https://fantasy.premierleague.com/>

weekly production will accrue to the team score and promote one of their squad members to be the captain (i.e promotion), who has his productivity doubled. The remaining four players are substitutes, whose weekly production will only count towards the firm's production if a nominated squad member does not enter the pitch in that week. We have information on hiring, firing and captain selection on a weekly basis allowing us to track the dynamics of discrimination within firms and thus test the predictions of Becker's taste-base model.

Our analyses examine what role, if any, workers' race plays in employers' choice of their squad at the start of the season and their subsequent decisions to recruit, retain and promote footballers each week of the season, conditional on their productivity, price and other factors. This allows us to assess whether discriminatory behaviors change over time, as the competitive environment changes (Levitt, 2004).

In an earlier paper, using aggregate weekly data on hires and fires we found little evidence of racial discrimination in the FPL (Bryson and Chevalier, 2015). However, our aggregate week-by-week data meant we were unable to observe team-level dynamics, so we were unable to establish the marginal productivity of white and non-white players, something which is critical in testing Becker's (1957) model which points to the marginal discriminator as crucial in establishing the degree to which non-white players need to be more productive than their white counterparts on the team to be hired, selected into the squad, promoted or avoid being fired. With these micro-data on FPL participants, we find that employers discriminate on racial grounds when deciding who to hire, who to select to a point paying position, who to promote to team captain and who to fire. Consistent with Becker's model of taste-based discrimination, non-white players need to be more

productive than white players to be recruited in the first place, to avoid dismissal, and to become team captain. This behaviour is stable over time, and is unaffected by the competitive environment, the geographical origin of the participant or their gender.

Furthermore, since we are able to measure the weekly team's production, we are in the unique position to measure the costs of discrimination: the distance between weekly team production and the maximum achievable in that week. We find that discriminating firms only suffer from a small loss of productivity since professional footballers are close substitutes.

Finally, we replicate these results using three different metrics to establish ethnicity based on anthropomorphic measurement, name or skin tone.

## **2. PREVIOUS LITERATURE**

Taste-based discrimination occurs when employers' decisions to offer jobs to applicants, or set wages, are based on prejudice with respect to specific traits rather than aptitude for the job (Becker, 1957). Discriminating firms may suffer from reduced productivity if their recruitment or promotion procedures are based on this prejudice rather than aptitude. Becker also noted the possibility of discrimination stemming from customers' and employees' taste preferences. Rather than taste-based discrimination, Arrow (1972, 1973) and Phelps (1972) emphasise asymmetry of information as a source of discrimination whereby, in the absence of information on prospective employees' productivity, employers judge the quality of applicants based on group characteristics such as race, resulting in potentially discriminatory outcomes. As noted by Levitt (2004: 433)

“in general, empirical tests have a difficult time distinguishing between taste-based and information-based models of discrimination”.

Causal evidence of discrimination in the labour market mostly relies on correspondence studies, or experimental set-ups correspondence studies identify employer discrimination in hiring by submitting made-up curricula vitae differing only by implied race to real job vacancies (see Bertrand and Mullanaithan, 2004, for a typical example, and Baert et al, 2018, for a review).<sup>2</sup> These studies, replicated around the world usually find a reduced likelihood of an employer “call back” for black applicants relative to white but otherwise identical candidates. Neumark (2018) offers a review of the experimental evidence on discrimination, concluding that the evidence on hiring discrimination is “most evident with respect to race and ethnicity” (p. 857). Audit studies are clean in the sense that they are able to isolate the role of race on hiring through the manipulation of curricula vitae, but they suffer from the fact that no actual hiring takes place. What they observe instead are 'call-backs' or offers; i.e. discrimination at an early stage in the recruitment process.<sup>3</sup>

Other types of study have attempted to differentiate between taste-base and statistical discrimination. Using data for the United States on racial attitudes by state linked to the Current Population Survey (CPS) Charles and Guryan (2008) support Becker’s prediction, stated above, that relative black wages are set by the marginal discriminating employer. They estimate that around one-quarter of the racial wage gap is due to taste-based discrimination. Relying on a dynamic set-up, Altonji and Pierret (2001) start from

---

<sup>2</sup> See Bendick (2007) for a review of audit studies providing evidence of employment discrimination on grounds of race. Stijn Baert maintains a repository of audit studies on discrimination: [http://users.ugent.be/~sbaert/research\\_register.htm](http://users.ugent.be/~sbaert/research_register.htm)

<sup>3</sup> Cahuc et al. (2019) argue discrepancies between the discriminatory attitudes of French employers towards North Africans revealed in surveys are underestimated in correspondence studies because discrimination at the invitation stage is a poor predictor of discrimination at the hiring stage.

the premise that, if employers discriminate on statistical grounds, wages will become less (more) correlated with easy-(hard)-to-observe worker characteristics over time as employers learn about their workers' productivity. They show that employers do appear to discriminate statistically with respect to easily observed characteristics, but not in the case of race, leading them to conclude "our estimates suggest that statistical discrimination on the basis of race plays a relatively minor role in the race gap in wages" (p. 316). Bohren et al. (2018) rely on prior evaluations of a posted question on a forum to assess potential biases in group-based belief about ability. In a similar way to Altonji and Pierret (2001) they observe belief reversal over-time. In their experimental set-up women face initial discrimination but are later perceived as more reputable than men; which is consistent with their model of biased-belief (i.e. stereotype) discrimination.

A distinct source of discrimination might be attention deficit. In a laboratory experiment Dovidio and Gaertner (2000) find employers discriminate on racial grounds, but only in the case of applicants whose qualifications mean the hiring decision is a difficult one. Charness and Kuhn (2011) summarize the experimental literature on discrimination. In the laboratory, racial discrimination appears to stem mostly from stereotyping. In a field experiment, Bartoš et al. (2016) link discriminatory behaviour to endogenous allocation of attention, whereby foreign names reduce effort to examine curricula vitae. Our set-up will also allow us to shed light on this source of discrimination.

In the sports context, the consensus is that racial discrimination has declined over time. Reviewing the wage discrimination literature for the United States, Rosen and Sanderson (2001: F58) suggest that discrimination "was easily detected in the initial studies of the 1960s and 1970s [But] It is difficult to find a negative coefficient on race in US data

these days". However, Altonji and Blank (1999: 3196) argue that there is evidence of salary discrimination, especially in professional basketball, some customer discrimination against minority players, and "some hiring discrimination, although these results depend on the sport and position [of the player on the field]".

Racial discrimination may be less apparent than it used to be because black players have been integrated into North American professional sports. Goff et al. (2002) treat the integration of black players into North American baseball and basketball as akin to the diffusion of a productivity-enhancing technology. Consistent with this proposition they show black players were more productive than white players during the quarter century over which sports moved from a segregated to an integrated equilibrium. The productivity differential dissipates post-diffusion. Szymanski (2000) finds similar results for the English professional soccer. Bryson and Chevalier (2015) use data on the aggregate demand for footballers in the fantasy football league, a set-up which is only consistent with statistical discrimination and report that conditional on performance, there is no effect of race on labour demand. This is, however, uninformative about the behaviour of the marginal discriminator which is the focus of this paper.

### **3. INSTITUTIONAL SETTING AND DATA**

#### *3.1 Institutional Set-up*

We analyse the virtual market of the Fantasy Premier League (FPL), an on-line game based on the Barclays Premier League, the top professional football league in England. FPL participants manage a virtual team of real professional footballers.



Participation in the FPL is free, there is no money involved and only the managers of the three teams with the most points at the end of the season are rewarded with a VIP trip to a premier League game of their choice. Monthly rewards are also available.

FPL participants receive a fictional budget of £100 million from which they must purchase a team of fifteen among the 670 professional footballers playing in the Premier League.<sup>4</sup> Like in a real firm, different positions must be filled; a team consists of two goalkeepers, five defenders, five midfielders and three strikers. There is no restriction on how many FPL participants can “own” the same footballer, as such hiring decisions are not affected by strategic decisions and purely reflect preference for the footballer.

The virtual team produces fantasy points based on the realised actions of the selected real footballers on the pitch. Footballers score points for playing in a particular week, for the time spent on the pitch, and for the actions they perform (positive points for goal scoring, assists and the like, and negative points for own goals, and disciplinary offences leading to red and yellow cards) and bonus points for overall performance. Bonus points are awarded to the best three players in each game, again based on some pre-determined metrics (see Appendix). Productivity is thus objectively measured for all workers and potential workers at no cost.<sup>5</sup> The role of the manager to maximise this production by selecting footballers subject to a budget constraint. The manager with the most points at the end of the season is the winner.

---

<sup>4</sup> They cannot select more than three footballers playing for the same club in the Premier League thus limiting participants’ opportunities to indulge preferences for players from teams they support.

<sup>5</sup> The productivity measure is largely independent of the referee – apart from disciplinary offences - so the behaviour of players is unlikely to be directly affected by referees’ discrimination (Parsons, Sulaeman, Yates and Hamermesh, 2011).

The Premier League is played over 38 periods (20 teams playing each other twice). Participants can hire and fire new workers after each game, subject to budget constraints. The price of players is set by the FPL and reviewed every week. The cost of a hire is the value of the incoming player plus the gap between the value of the outgoing player on the open market and the value the employer recovers on sale (which is not the full market price).<sup>6</sup> Employers are permitted one transfer per week at no fee. Any additional transfers in a week entail a deduction of four points, which must be added to the financial cost of making a transfer.<sup>7</sup> As well as squad selection, the manager selects the 11 point-scoring footballers (the team) whose production will count towards the manager's performance. The remaining four players are ranked, and might become eligible to score for the manager if one of the team member does not play in that week. The manager also appoints a captain whose productivity in that week is doubled. A vice-captain is also selected, which will take the position of captain in case the captain does not play in that week.

This setting has several features that allow us to identify the effects of taste-based employer discrimination as opposed to other types of discrimination. This is because i) employers have perfect, up-to date information on the productivity and price of all workers in the industry; ii) production is simply additive in each workers' production and there are no productivity spill-over across workers; iii) firms do not have customers ; iv) workers have no say in the firms they join/exit<sup>8</sup>; v) workers are able to work at more than one firm

---

<sup>6</sup> There is a gap between the buying and selling prices of players. This margin is half of the difference between the current price and the price at which the player was bought; this can be thought of as a tax on the value added. As such, transferring players has a financial cost and leads to a reduction in the firm's budget. So firms may not always optimize their teams and may refrain from using their weekly transfer.

<sup>7</sup> Once per season FPL participants are allowed to hire an unlimited number of footballers with no point penalty.

<sup>8</sup> Worker selection based on perceptions of discriminatory tendencies in particular occupations or among certain employers may contribute to labour market segmentation and wage discrimination. For example, Plug et al. (2014) find gays and lesbians in Australia shy away from more prejudiced occupations.

simultaneously - as such hiring reflects true preferences, subject to firms' budget constraints, and is not affected by strategic concerns. At the beginning of the season, the information on player's productivity refers the past season, if available, or is missing for footballers new to the premier league, it is thus possible to argue that some statistical discrimination might be possible in the first few weeks.

Second, being a virtual labour market, participants to the FPL are free to discriminate when hiring or firing. In this sense we are "turning back the clock" to a time when employers faced no legal impediments to discrimination. Therefore, the costs of discrimination are low and we can thus identify an unbiased taste for discrimination.

Third, employers are price takers: the price of recruiting individual workers, i.e. a sign-on fee, varies substantially but individual employers are unable to influence these prices. Prices attached to workers are exogenous to the firm, but relate very strongly to worker performance, as we shall see. As such, firms cannot exploit minority workers by offering them lower sign on fees. Once signed, there are no wages in our set up. Therefore discrimination can only be expressed in terms of hiring, firing and promotion probabilities.

Fourth, ours is a dynamic setting with a large number of employers engaged in many transactions over 38 periods. We are therefore able to estimate influences on employer hiring and firing decisions over time. There is no left censoring in the hiring history since all participants must hire their squad at the start of the season. (Participants may of course have a history of playing the game but we are able to account for this). We can thus test whether discrimination changes over time. There are several reasons to believe

---

Similarly, co-workers might sort according to their racial preferences: Hedergaard and Tyron (2018) use a field experiment to show that, consistent with taste-based discrimination, co-workers are willing to forego 8 percent of earnings to work with people from the same ethnic group.

that discrimination may change over time. Participants might learn about the losses resulting from discrimination over time, the costs of discrimination change over the course of the season and, as the season progresses, the signal to noise ratio of the productivity history improves eliminating the scope for statistical discrimination.

One might question whether employers would express their discriminatory preferences in a set-up like the FPL since there is no physical interaction between employers and employees.<sup>9</sup> However, the need for physical interaction to discriminate is not supported by psychological studies, and some of the physiological reactions associated with interactions with out-group members can be triggered with a mere picture. For example, Hart et al. (2000) report increased amygdala activities, a subcortical structure of the brain involved in perceived danger, when people are shown pictures of “out-group” individuals. Indeed, discrimination in on-line market is observed, even in set-up with no direct interactions between buyers and sellers (Laouenan and Rathelot, 2022, among others).

A second concern is whether FPL participants will spend the time and effort to optimise in what seems to be a low-stakes game. A lack of effort might be associated with more discrimination if participants rely on some heuristics rather than actual productivity measure when taking their decisions. However, anecdotal evidence suggests that participants spend a considerable amount of time thinking about their team picks and adjusting their squad, and that peer pressure rather than financial gains is an important motivator.<sup>10</sup> Peer pressure can be approximated by performance in local leagues. Local

---

<sup>9</sup> Alternatively, one could argue that the absence of interaction makes it easier to discriminate.

<sup>10</sup> When interviewed by the media Simon March, the 2015 FPL winner, said he was constantly thinking about how to improve the performance of his team (BBC 04/08/2015). A 2008 survey of fantasy league

leagues are self-selected group of participants whose performance and local ranking is reported to the selected participants. We will later explore the effect of this local competition.

### *3.2: Player Data Description and inferring race*

Not all of the footballers listed in the FPL are available for the maximum 38 games in the season. The premier league starts before the trading market closes, so the set of footballers expand as clubs finalised their teams.<sup>11</sup> In mid-season clubs are allowed for a short period of time to trade players, and towards the end of the season, young players might become eligible to join the main squad. 670 footballers are listed for at least one week in the FPL, leading to an unbalanced panel of 23,619 weekly observations. The information collected is the footballers' weekly production of FPL points, market value, as well as the number of net transfers transfer for this player in FPL and the share of FPL managers who "own" this player. To this we add, non-varying information on the club, nationality and whether the footballer has ever played for his national team.

The race of footballers is inferred in three different ways. When accessing the complete information set on the performance of a footballer, participants see a picture of this player. We use this picture to define ethnicity from i) an AI-powered software that measures anthropometric features and ii) by rating skin tone. Note that participants might discriminate based on players' name, in which case they will not see the picture but will

---

participants in the United States revealed that the three main reasons to participate were enjoyment, competition with others and pressure from peers (Baerg, 2009).

<sup>11</sup> Some footballers listed in week 1 are transferred out of the premier league, and while still listed on the FFL website, their performance is no longer monitored; since these footballers remain available to FFL participants, and are indeed selected by some participants, we keep them in our dataset but include a dummy variable to indicate footballers no longer active in the premier league.

also forfeit the more detailed information on the player's performance. We thus also approximate race by the name of the footballer, using an algorithm provided by NamSor<sup>12</sup> (see Appendix for detail). From these measures we compute dichotomous indicators of race. According to these three measures, the proportion of non-white footballer is between 0.28 and 0.34. The AI and skin-tone measures have a correlation of 0.83 but the name-based indicator only has a correlation of 0.40 with the other two measures.

Footballers in FF are given a fixed position on the pitch, which reflects their position on the pitch and affects their productivity (see section A1.1.2). There are four such occupations: goalkeeper, defender, midfielder and forward. The ethnic composition differs substantially by position (Table 1). Using the AI measure, only 4% of goalkeepers are non-white but 41% of forwards are. Very similar splits are found with the skin-tone measure, while the name measure has less variance: 11% non-white goalkeepers and 32% non-white forwards. For all ethnic measures, non-whites are significantly under-represented as goalkeepers. Non-whites are also less likely to be UK players but the other fixed characteristics are very similar between players of different race, even if this differs by the race definition (Table 1- Panel A).

### *3.3 Footballers productivity and Race*

We now assess to what extent there are any racial differences in the distribution of productivity between white and non-white footballers. We use the AI measure here, but evidence for the other measures of race are available in appendix. The first panel in Figure 1 reports a large overlap in the distribution of the total number of points scored over the season by race, the average total points accumulated over the season by footballers are not

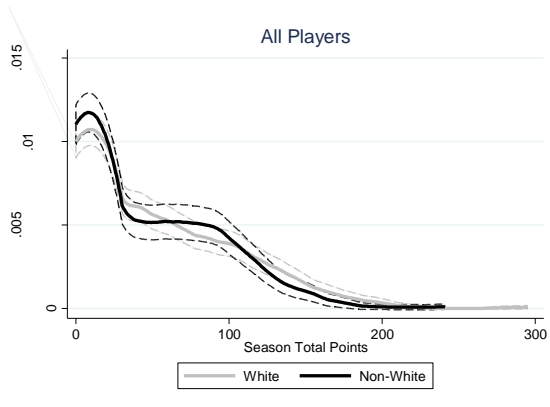
---

<sup>12</sup> <https://www.namsor.com/>


significantly different between race (48.9 vs 42.9). The remaining panels in Figure 1 pertain to variations in how these points might be accumulated. However, we do not find any difference in the number of minutes played, standard deviation in total points or the period of the season in which the points are accumulated. Overall, as asserted by Goff et al. (2002) top-flight professional sports are racially integrated and professional footballers are largely substitutable by ethnicity. This is true, even within on-field position (see Appendix).

**Figure 1: Productivity by Race (AI)**

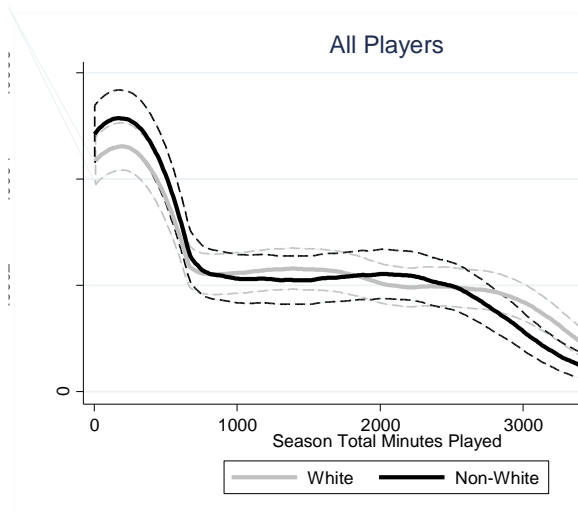
A] Total Number of points over the season



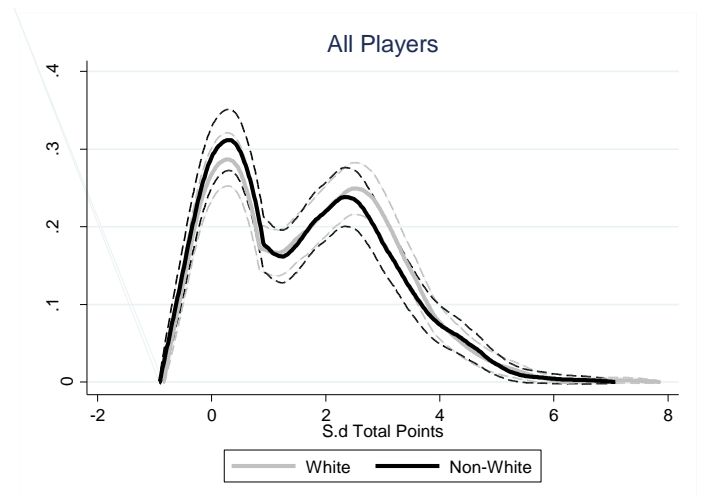
C] Points per week of the season



B] Total Minutes played over the season



D] Standard Deviation in Total Points



The second main determinant of demand for a player is price. Prices are set by the FPL and appear to be a function of performance and net demand (see below), and range between £3.8 million and £14 million. After the initial valuation, players' prices are reviewed by FPL on a weekly basis: we show in the next section that this is largely driven by performance and demand and is not systematically related to race. The second panel in Table 1 reports weekly varying variables by race, separately for the three different race measure. Not adjusting for performance, non-white players are less likely to have played at all, produce fewer points, and have higher values than white players, this makes them less likely in demand, which fewer teams own on average.

### *3.4 Does the FPL function like a labour market?*

Before we investigate the racial differences in hiring we need to establish whether the FPL functions like a labour market. In Figure 2 Panel A we explore the relationship between performance in the last period and current net aggregate demand for a footballer by FPL participants. We would expect a positive relationship if participants expect performance to be positively correlated between weeks (i.e. reflects ability and effort rather than luck). Indeed, the correlation between previous week performance and net aggregate demand is positive and does not significantly vary by ethnicity. This indicates that participants to FPL are putting effort in managing their team, consistent with trying to maximise output. At this level of aggregation and for the area of common support, we also find no differences by ethnicity in the demand for players, conditional on performance.



**Table 1 Descriptive Statistics (Player level)**

<b>Panel A: Fixed Characteristics</b>						
<b>Race definition:</b>	AI		Name		Skin	
	White	Non-white	White	Non-white	White	Non-white
Defender	0.344	0.268*	0.335	0.289	0.320	0.304
Forward	0.160	0.249**	0.168	0.226	0.171	0.230
Goalkeeper	0.152	0.033***	0.136	0.069**	0.150	0.042***
Midfielder	0.344	0.450**	0.361	0.415	0.359	0.424
UK	0.494	0.327***	0.582	0.255**	0.504	0.322***
EU15	0.312	0.279	0.259	0.343*	0.323	0.265
Other	0.110	0.011***	0.085	0.053	0.106	0.021***
European						
Americas	0.062	0.138***	0.031	0.160**	0.044	0.159***
Africa	0.000	0.204***	0.017	0.154**	0.005	0.187***
Others	0.022	0.041	0.026	0.035	0.018	0.046*
International	0.581	0.643	0.554	0.664**	0.576	0.647
Points <sup>+</sup>	62.531	54.670	61.427	56.748	62.291	55.396
In FF <sup>+</sup>	0.653	0.710	0.679	0.673	0.649	0.714
Years in FF	2.451	2.688	2.736	2.336*	2.426	2.710
Observations	401	269	352	318	387	283

<sup>+</sup>Last season

---


<b>Panel B: Weekly Characteristics</b>						
Minute	33.301	29.452**	33.31	30.050*	32.105	31.308
Played	0.443	0.414***	0.438	0.423*	0.431	0.431
Yellow card	0.053	0.048	0.050	0.053	0.052	0.050
Red card	0.002	0.003	0.002	0.002	0.002	0.003
Weekly points	1.384	1.259***	1.354	1.311	1.338	1.329
Point per min	0.059	0.059	0.057	0.062*	0.059	0.058
Points per £	0.255	0.224***	0.256	0.227**	0.245	0.238
Net Transfer	350.19	-175.257	34.56	258.828	180.374	87.052
Current value	5.039	5.260***	4.988	5.282**	5.057	5.222***
% Teams	2.693	1.991***	2.307	2.531**	2.625	2.126***
Observations	14207	9412	12430	11189	13605	10014

Note: \*\*\*, \*\*, \* indicates significant difference between the 2 groups at 99%, 95%, 90% level, respectively

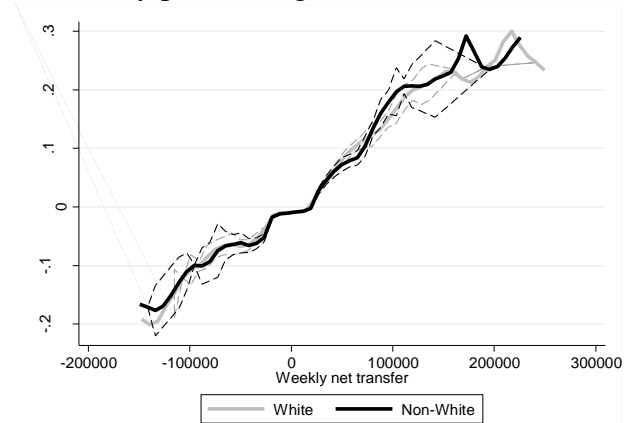
As in any competitive market increased aggregate demand correlates with higher prices (Panel B), indeed the two variables move in an almost linear way, with no difference by ethnicity. Overall, the FPL appears to behave like a labour market, where more productive workers are in higher demand and where demand and price are positively correlated, and these basic relationships do not differ by race.

**Figure 2: Does the FPL functions like a labour market**

A) Lagged player performance and net demand



B) Weekly price change and net demand

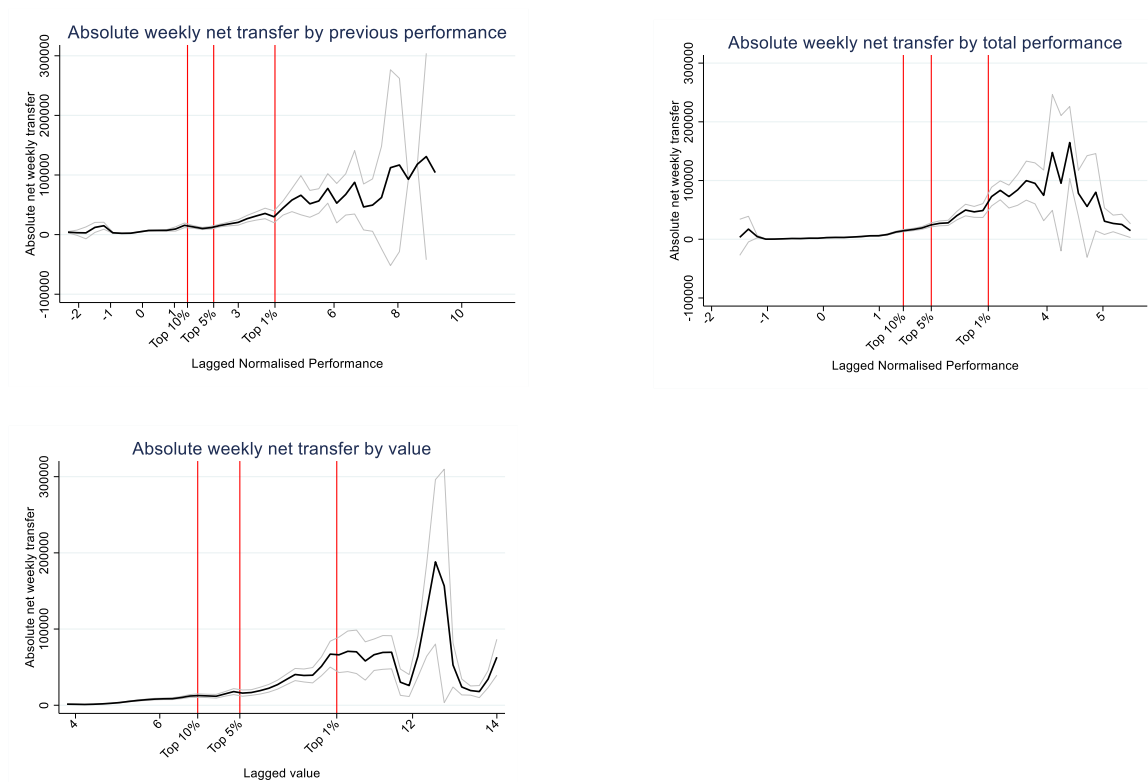


Note: The thick lines report local polynomial smooth plots separately by ethnicity. The thin dash lines are the 95% confidence interval around these plots. Weekly net transfer refers to the net demand for footballers, lagged performance is the footballer's productivity in the previous period, and weekly price variation is the change in price compared to the previous week. Graph A is based on 23,496 observations; i.e. 670 footballers observed for a maximum of 37 weeks (week 1 is excluded since there are no previous values of price or performance). Graph B is based on the full 38 weeks. Race is defined using the AI definition.

Since trading more than one player is costly – 4 production points for additional transfer – most participants if trading, only trade one player. This makes the decision particularly important. As such, most of the trading involve the best performing players. Figure 3 reports the absolute values of the net weekly transfers for different metrics of star performance (lagged production, lagged total production, and log value). For each graph,

the vertical lines highlight performance at the top 10%, 5% and 1%. The transfer activities is almost restricted to the top 5% performers.

**Figure 3: Absolute Net Weekly Transfers by Performance**



The FPL does not communicate how footballers values are fixed, but it appears that they are driven by relative demand. As such they could reflect the racial biases of participants. In Table 2, we report estimates on weekly productivity differential by race, accounting for footballers fixed characteristics. In column 2, we account for the player's lagged value, while in column 3 we include players fixed effects. In none of the specifications, do we find evidence of productivity differential by race.

**Table 2: Player's level Outcomes by Race – AI definition**

	Points	Points	Points	log. Value	Log. Value	Log. Value	% team	% team	% team
Non-white	-0.176 (0.104)	1.605* (0.707)		-0.019 (0.015)	-0.023 (0.013)		-0.752 (0.428)	1.387 (1.831)	
Lagged value		4.608*** (0.506)	-3.580** (1.202)					1.614*** (0.348)	8.512*** (1.031)
Lagged Total Points					0.003*** (0.000)	0.001*** (0.000)		0.089*** (0.010)	0.017* (0.007)
Non-white # Lagged value		-1.055* (0.452)	2.372 (1.603)					-0.248 (0.367)	-2.785* (1.135)
Non-white # Lagged Total					0.001 (0.000)	-0.000 (0.000)		-0.018 (0.012)	-0.011 (0.008)
Obs.	23619	22949	22947	23619	22949	22947	23619	22949	22947
Time fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Player fixed effect	No	No	Yes	No	No	Yes	No	No	Yes
Mean for White	1.384	1.384	1.384	1.593	1.592	1.384	2.693	2.703	2.704
Size effect s.e	-0.127 (0.075)	-0.065 (0.060)	. (.)	-0.012 (0.009)	-0.006 (0.008)	. (.)	-0.279 (0.159)	-0.124 (0.105)	. (.)

Note: OLS estimates. Additional control for club, region of origin, experience, position on the field, international and homegrown status, no longer playing in league, number of games played in the week, and weeks of play. Standard errors adjusted for clustering at the player level (670 clusters) are reported in parentheses. \*\*\*, \*\*, \* denote statistical significance on the 10, 5 and 1 level respectively.

In column 4 to 6 we assess the value of footballers. In none of the specification do we find differences in valuations of white and non-white footballers. Finally in Column 7 to 9, we assess the fraction of teams that have a given footballer in the team, and do not find differences in the popularity of footballers by race.

### 3.5 Participants to the FPL

The data available to us mainly pertain to the weekly squad and team selection of participants and captaincy/vice-captaincy decision.<sup>13</sup> This allows us to identify weekly hiring, firing and promotion decisions. The sample is restricted to 2,288,761 participants who engaged in the week from week 1 onwards, providing us with a balanced sample of 86,972,918 records over the 38 weeks season.

Since the availability of non-white players differs by position, we define the racial composition of the team based on the proportion of non-whites in different positions. More specifically, for each position we compare the numbers of non-white selected which sums the proportion of non-whites in different playing positions thus:

$$RI_{jt} = [\sum_p N_p * ((N_{pjt}^b / N_p) - \overline{N_p^b})] / N \quad (A1)$$

Where  $N_{pjt}^b$  is the number of non-white players in the  $p$ -position selected by team  $j$  at period  $t$ ,  $\overline{N_p^b}$  is the proportion of non-white players in position  $p$  available in the FPL, and  $N_p$  is the total number of players at position  $p$  in the team. The racial index is computed for the full squad (SRI) of 15 footballers selected, in which case  $N_p$  is fixed (2 goalkeepers,

---

<sup>13</sup> Note, we do not have information on which four players were originally selected to be part of the reserve; only on which were eventually left out; nor on whether the captain was chosen or automatically substituted with the vice-captain.

5 Defenders, 5 midfielders and 3 forwards); and for the weekly selection of 11 point-scoring squad members<sup>14</sup>; the team (TRI), in which case  $N_p$  can differ between teams.

By construction, a participant whose squad/team reflects the racial composition of the FPL would have an *ID* equal to 0. A negative index sign indicates a team that has more white players than would be expected if players were picked at random. A positive value, indicates a team with more non-white players than expected.

Table 3 reports the average value of these indices for each measure of race. The mean values of the Squad Racial Index over the full season are -8.06, -5.04, -3.70 for the AI, Name and Skin-tone measures respectively, indicating that on-average squads are composed of more white players than if picked at random.

We also define non-white discriminators as FF managers whose squad/team has a value of the Racial Index that is lower than minus one standard deviation of the Racial Index. Similarly, White discriminators are FF participants whose squads has a Racial Index greater than one standard deviation. Depending on the definition of race, 30% to 52% of managers have selected a squad that over-represents white footballers, and xx% to xx% one that over-represented non-white footballers.

**Table 3: Indices of Racial Composition**

	AI measure	Name measure	Skin-tone measure
Squad Racial index	-8.06 (11.72)	-5.04 (10.93)	-3.70 (12.44)
Non-White discriminator:	52.18 (49.95)	36.39 (48.11)	30.45 (46.02)
Squad			

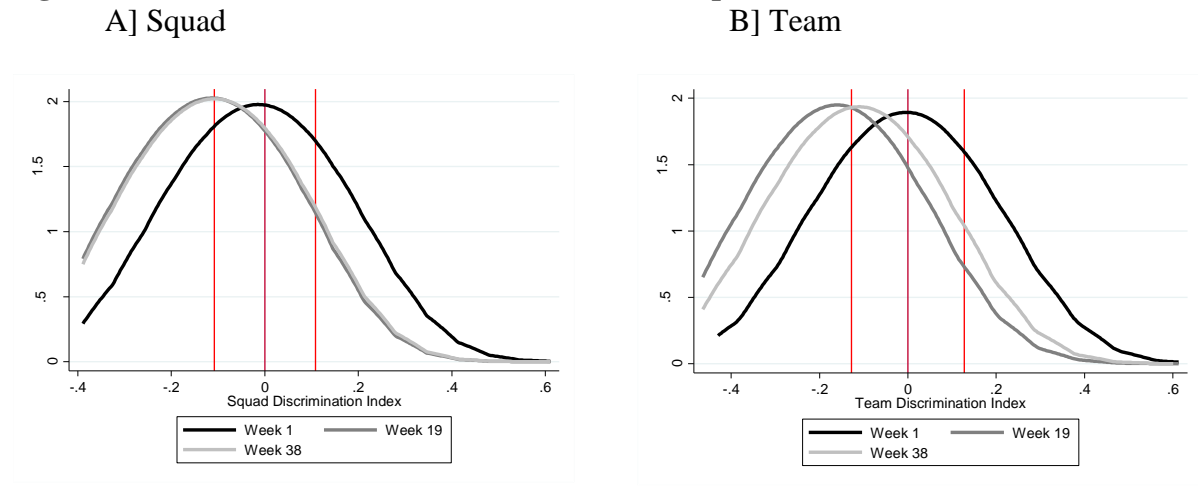
Note: Racial index is defined as a position weighted discrimination measure of squad/team, as in (A1). White discriminator is an indicator equal to one when the racial index has a value lower than 0 minus 1 standard deviation in racial index.

Sample: N=3,737,604 (98,358 \*38)

<sup>14</sup> A team must include one goalkeeper, at least one forward, 3 defenders and 3 midfielders, and at most 5 defenders, 5 midfielders and 3 forwards.

In Figure 3, we report the distribution of this index at the squad and team level for three specific weeks (Week 1, 19 and 38). The red vertical bars define a window of plus and minus one standard deviation in the index, compared to 0. For the AI measures the distribution shifts to the left (i.e. the share of white players increase, by the mid-season point (week 19) but remains stable thereafter. The evolution of the discrimination index at the team level is very similar. Similar Figures are available in the Appendix. They reveal that the evolution of discrimination differs with the definition of race used. There is little variation in the racial composition of teams when using the skin-tone measure, and when using the name-based measure, we find that if anything teams become less white over-time.

**Figure 3: Distribution of Discrimination Index in Specific Weeks (AI measure)**



Note: Team discrimination Index, based on the selected 11 point-scoring footballers, for the panel of managers still actively managing their team after week 34.

We also collect information on fixed characteristics of participants including their past experience in the FPL (number of years of participation, best rank, played in the

previous season), and their stated country of origin. These fixed characteristics of participants are reported in Table 4 for three different subsamples.

**Table 4: Fixed Characteristics of Managers by sub-population**

	Balanced Panel	Named Panel	Persistent Panel
Nbr of past Participations	1.861 (2.031)	1.925 (2.058)	2.383 (2.116)
Best previous rank	755961.0 (692710.3)	738764.3 (685226.2)	573940.4 (607439.0)
Participate in at least 1 local league	0.828 (0.378)	0.826 (0.379)	0.917 (0.276)
Size of local league	145.2 (738.8)	28.03 (73.35)	166.7 (837.2)
Nbr of local league	1.874 (1.799)	1.879 (1.797)	2.361 (2.006)
Last week of Activity	23.56 (14.70)	24.12 (14.58)	37.12 (1.197)
Non-White Participant		0.398 (0.490)	
Male participant		0.915 (0.279)	
UK	52.74	53.93	52.47
EU15	11.96	11.61	12.13
Other European	5.37	4.87	5.50
Americas	6.35	6.03	5.27
Africa	4.50	4.51	4.89
Asia	18.44	18.40	19.09
N.A.	0.64	0.65	0.65
Observations	2,288,761	98,358	1,002,201

Note: Balance panel of participants observed from week 1 to 38. Named participants, random sample of participants from the balanced panel for which race and gender was predicted from their self-declared name. Persistent panel, sub-sample of participants who engaging in trading activities at least once in the last 5 periods of the season (week 34 and above).



For a random-subsample of 98,358 participants we collect the stated names of the participants when registering to FPL. These names are used to predict gender and ethnic origin (see Appendix for detail). These participants define the named subsample. The football season last for 9 months, and one worry is that participants interest in the game vanes over time, and that the sample of participants who engage in the game, and thus contribute to the identification, becomes more and more selected over time. Indeed, after the first 5 weeks less than 50% of participants are engaging in transfer activities in a given week and by week 34, less than 50% of the initial participants are still actively involved in the game; we name these participants the persistent participants (see Appendix).

Table 4 reports descriptive statistics for the different sub-samples. The samples are similar in the origin of the participants. Slightly more than 50% declare that they live in the UK, Asian and the EU15 represent 18% and 12% respectively, and the remaining continents about 5% each. Participants have on average played twice previously and 80% of them are participating in at least one local league. On average, the last transfer participation was in week 24. Due to randomisation, the name sample has almost identical characteristics. Participants to Fantasy Football are 92% male and 40% are non-white. Participants in the persistent panel have slightly different characteristics.. They have 0.5 more years of experience playing FPL, have performed better in the past , and are more likely to participate in more and larger local leagues. They on average traded players until the week 37.

Each week, FPL provide the player selection that would have achieved the maximum amount of points. We use this measure as the weekly production frontier and compute the loss function as the difference between the frontier and the points achieved,

which we expressed as a percentage of the frontier; i.e. the greater the distance the less efficient the manager is. This distance will allow us to assess the costs of discrimination. We also collect the total number of points accumulated so far by each manager. A crucial element in Becker’s model is the competition that discriminators face. We approximate this with performance in the local league. In particular we compute, its (percentile) rank in the local league based on total points accumulated where - to account for differences in local league size - rank is defined as  $((rank_{li} - 1)/(N_l - 1)) * 100$ , where  $rank_{li}$  is the rank of individual  $i$  in local league  $l$  and  $N_l$  is the number of participants in local league  $l$ ; thus 100 is the best performer in the local league and 0 the worst. We also compute the point gap, in percentages, to the managers above and below in the local league ranks, as a proxy for competitive pressure. Finally, we compute the “leave-out mean” of the discrimination indices within a local league to approximate the social norm of the local competitors. These weekly varying variables are reported for the three sub-sample of interest- in Table 5.

#### **4- Modelling considerations**

The FPL is an ideal setting to test Becker’s model. Under Becker’s model of taste-based discrimination, the dis-amenity value in employing minority workers results in minority workers having to “compensate” employers by being more productive at a given wage or, equivalently, accepting lower wages for identical productivity. Employers maximise a utility function for profits (or, in our case, points won) subject to the monetary value of utility from employing members of particular groups, thus:

$$U_e = pF(Nb + Na) - w_a Na - w_b Nb - dNb, \quad (1)$$

where  $p$  is price at which production is sold,  $F$  is the production function,  $N_x$  is the number of workers in group  $x = \{a, b\}$  and  $w_x$  is the wage paid to members of the majority and minority groups respectively.  $d$  is the taste parameter of the firm (which Becker refers to as the “coefficient of discrimination”). Prejudiced employers ( $d > 0$ ) act as if the wage of  $b$  group members is  $w_b + d$  so that, assuming equal productivity between the two groups, they will only hire  $b$  group members if  $w_a - w_b \geq d$

**Table 5: Weekly Varying Manager Level Variables**

	Balanced Panel	Named Panel	Persistent Panel
Normalized Weekly points	0.000 (1.000)	0.0139 (0.999)	0.366 (0.901)
Normalized Total points	0.000 (1.000)	0.0224 (1.001)	0.595 (0.740)
Productivity loss (in %)	63.34 (12.45)	63.17 (12.46)	59.03 (11.91)
Team Racial Index (AI)	-0.0755 (0.139)	-0.0765 (0.139)	-0.108 (0.128)
Percentile rank in Local league	50.00 (33.42)	50.63 (33.50)	64.27 (30.41)
Gap to person above in league	4.785 (109.4)	4.730 (46.05)	3.314 (19.13)
Gap to person below in league	-3.903 (5.806)	-4.014 (5.833)	-4.090 (5.663)
Mean racial index in Local League (AI)	-0.0636 (0.0564)	-0.0638 (0.0567)	-0.0718 (0.0553)
Observations	86,972,918	3,737,604	38,083,638

In the FPL wages are set exogenously beyond the firm, so employers can only exercise their taste for discrimination in hiring, firing and promotion decisions. By analogy, and assuming race-specific production functions  $F_a$  and  $F_b$  respectively, in an environment where wage discrimination is impossible ( $w_a = w_b$ ), the last  $b$ -type worker is hired only if  $F_b - d > F_a$ . Where employers are unable to adjust wages they will only employ minority workers at the margin where the employer benefit from the perceived additional productivity of a  $b$ -worker ( $F_b - d$ ) offsets the utility derived from employing (retaining/promoting) an  $a$ -employee ( $F_a$ ).

First we provide some descriptive evidence that these predictions are indeed realised in the virtual labour market we assess. Almost one-third (30%) of FPL professional footballers are non-white, but they only constitute 22% of those selected in a team. Non-whites are also over-represented in the population fired from employment (27%) and under-represented among those promoted to be captain (17%).

We test four propositions arising from Becker's taste-based discrimination model. First, lower perceived productivity of otherwise identical  $b$ -workers relative to  $a$ -workers means  $b$ -workers have a lower (higher) probability of being hired (fired) compared with  $a$ -workers at a given level of productivity. Second, for the same reason, at a given level of productivity  $b$ -workers have a lower probability of being promoted to captain than an equally productive  $a$ -worker. Third, the marginal  $b$ -worker will be more productive than the marginal  $a$ -worker on the team. Fourth, the differential hiring and firing of  $b$ -workers compared with  $a$ -workers results in a compositional shift towards  $a$ -workers during the course of the season.

Table 6 provides supporting evidence on the productivity of the marginal workers which is consistent with Becker’s prediction. Defining the marginal worker as the least productive worker of each type we find that in teams with less b-type workers than if randomly selected, the marginal b-type worker is more productive than the marginal a-type worker. Indeed, the more of one type of worker the lower the productivity of the marginal worker of that type. We also note that the average differences between *a*- and *b*-workers are relatively small, and indeed non-existent for the non-discriminating team (3 *b*-worker), but are quite large for the marginal workers. The literature that estimates discriminatory behaviour based on the average worker might thus considerably under-estimate the costs of discriminating.

**TABLE 6: Average and marginal productivities**

Nb Non-White in team	Non-White Player		White Player	
	Average	Least Productive	Average	Least Productive
0			4	0.4
1	4.4	4.4	4.1	0.4
2	4.3	2.3	4.2	0.5
3	4.2	1.5	4.2	0.5
4	4.0	1	4.1	0.6
5	3.8	0.7	4.1	0.6
6	3.6	0.5	4.0	0.7
7	3.4	0.3	3.8	0.9
8	3.3	0.3	3.8	1.2
9	3.4	0.2	3.4	1.8
10	3.5	0.2	3.8	3.8

In the final column, we provide a measure of the loss associated with discrimination by computing the ratio of the absolute difference in the productivity of the marginal workers from both groups over the total production of the team in that week. The marginal productivity losses for teams with only 1 or 2 non-white players, are 8 and 4 percent respectively. Teams with more than half black players are also losing between 2 and 10 percent of productivity from their marginal worker. Table 6 thus stresses the importance of considering the marginal and not the average worker when investigating discrimination.

In order to identify the decisions made by participants we create all dyads of participants and footballers. Our first set of estimates focus on hires made to the squad at the start of the season, what we term first picks. At that moment, employers have less up-to-date information on player productivity than they have subsequently in the season. We run linear estimates for the probability of player  $i$  being picked by manager  $j$  for the first game of the new season  $H_{ijt=1}$ . We run four model specifications beginning with the simple raw correlation with being a non-white player  $b_i$ , then adding player characteristics  $X_i$  including price and past performance in the FPL,  $P_{it-1}$ . For players new to the Premier League there is no lagged FPL performance. We identify these players with a dummy variable. In the final model, we interact ethnicity and player productivity and with price as in equation (2), to assess whether this information is processed differentially, in a way that would reflect discriminatory behaviour:

$$H_{ijt=1} = \alpha + \beta_1 b_i + \delta X_i + \gamma z_j + \gamma P_{it-1} + \beta_2 b_i * P_{it-1} + \varepsilon_{1i} \quad (2)$$

Note, we also control for the participant  $j$  observable characteristics  $\mathbf{z}_j$  including, predicted gender and ethnicity, as well as self-reported region of origin, past participation in FPL, performance in past participation (last year and best ever), and numbers of private leagues entered as a proxy for interest in the game. For those with past experience in FPL, we also include their best rank to approximate their managerial competencies. The standard errors are clustered at the participant level to account for possible correlations in the decisions taken by a participant.

We treat the first period separately since discriminatory behaviour might be altered by increased effort in team selection – this is the only time where the 15 players can be selected at no cost – and less reliable information regarding players’ productivity; 15% are new to the premier league while others might have transferred to different clubs or have experienced substantial changes in their team mates or club manager. For the remaining 37 periods, we estimate the probability of player  $i$  being selected by team  $j$   $H_{ijt}$  conditional on having been a team member in the previous period  $H_{ijt-1}$ , some measures of price ( $Pr_{it-1}$ ) and performance in the last week ( $P_{it-1}$ ), and in the final model their interactions with race. The model also includes employer fixed effects ( $\gamma_j$ ). The estimated model is thus:

$$P(H_{ijt} = 1) = a + b_1 Black_p + d_2 H_{ijt-1} + d_3 Black_p * H_{ijt-1} + \sum_{k=1}^3 b_{2k} P_{it-k} + \sum_{k=1}^3 b_{3k} Black_p * P_{it-k} + c_2 Pr_{it-1} + c_3 Black_p * Pr_{it-1} + \gamma_i + \gamma_t + \varepsilon_{ijt} \quad (3)$$

As Kahn (2009: 14-15) notes, identification of racial bias in hiring and firing decisions is best investigated using performance differences of marginal workers, as opposed to the average worker because only the former is informative about the margin

where the hiring/firing decision is made. This is precisely what we observe and (3) can be interpreted as the probability of a team member being fired, and whether ethnicity alters this probability.

Whereas at the start of the season employers must rely on player performance information from the previous season, employer information about player performance is continually up-dated throughout the season such that, if there is any uncertainty regarding productivity at the outset that could induce some statistical discrimination, this dissipates over time as employers observe players' "form", including that of footballers they have not selected.

Another unique attribute of our dataset is that we can focus on another set of marginal players, the team captains. Each week, participants nominate one of their players as the team captain. The points accruing to this player in that week are doubled.<sup>15</sup> The costs of discriminating on captaincy are twice as large as on any other players. For this decision, we re-estimate a model akin to (2) for week 1 and (3) for subsequent week, substituting the hire variable for one indicating that the player has been chosen as the designated captain for team  $j$  in week  $t$ .

FAs seen above, the descriptive statistics support one of the main conclusions of the Becker model of discrimination that discriminating firms are less productive; i.e. their marginal worker is of lower productivity than in a non-discriminating firm. We further test this hypothesis here using the detailed information on the weekly productivity of firms. Moving back to participant-level data, we estimate the We then estimate the costs of

---

<sup>15</sup> The participant also nominates a vice-captain. If the nominated captain does not play that week the vice-captain's points tally is doubled. The data available to us refers only to the player whose productivity was doubled but we cannot identify whether this was the initial nominated captain or vice captain.



discriminating as the following linear model in equation (7), allowing for the effect of discrimination to differ between positive and negative values of the discrimination index.

$$LP_{jt} = \alpha_1 + \alpha_2 I(ID_{jt} < 0) * ID_{jt} + \alpha_3 I(ID_{jt} \geq 0) * ID_{jt} + \beta X_{jt} + \gamma_j + \gamma_t + u_{jt} \quad (7)$$

Since the reasons to discriminate might change over time, we also estimate whether the effect of discrimination on the loss function are time dependent.

## 5- Results

### 5.1: Initial decision

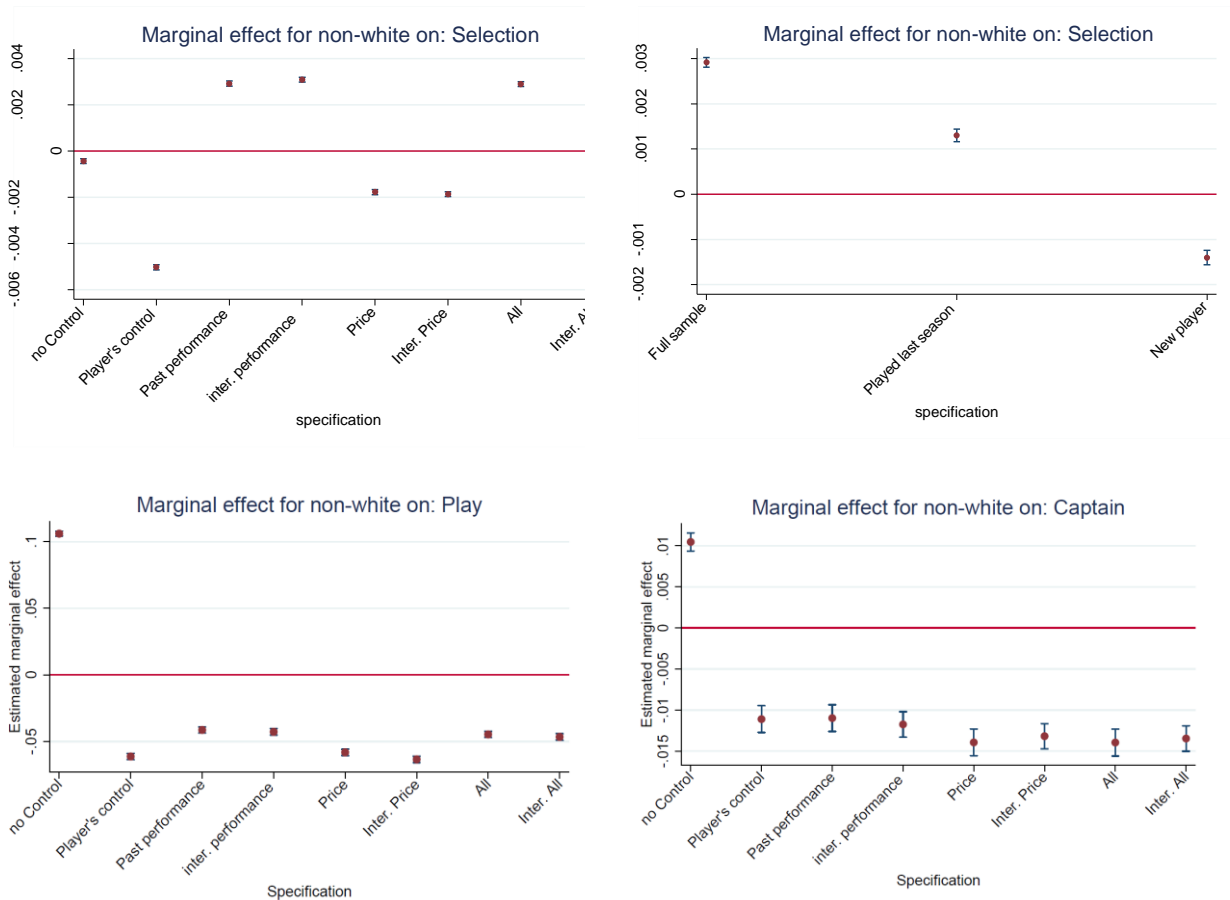
Figure 7 summarises the finding of estimating the model of squad selection in week 1 for different specifications. With the basic set of controls available in most studies, we find that Non-White players are 15% less likely to be selected in a squad. However, when accounting for past performance, the effect of race flips, and non-white become more likely to be selected. This positive effect of race remains when allowing the effect of past performance to differ by race. Replacing past performance by the value of player, results in non-white being discriminated; but including both performance and value, and their interactions with race leads to non-white being about 5% more likely to be selected in a squad in week 1.

In week 1, the measure of productivity is at best dated and absent for players who are new to the Premier League. Separating by players' past presence in the Premier League, (Panel B), we find that non-white players new to the league are indeed discriminated.

Focusing on the decisions of selecting the point-accruing players and the captain (Figure 8) reveal the importance of focusing on marginal decisions, While the evidence of racial

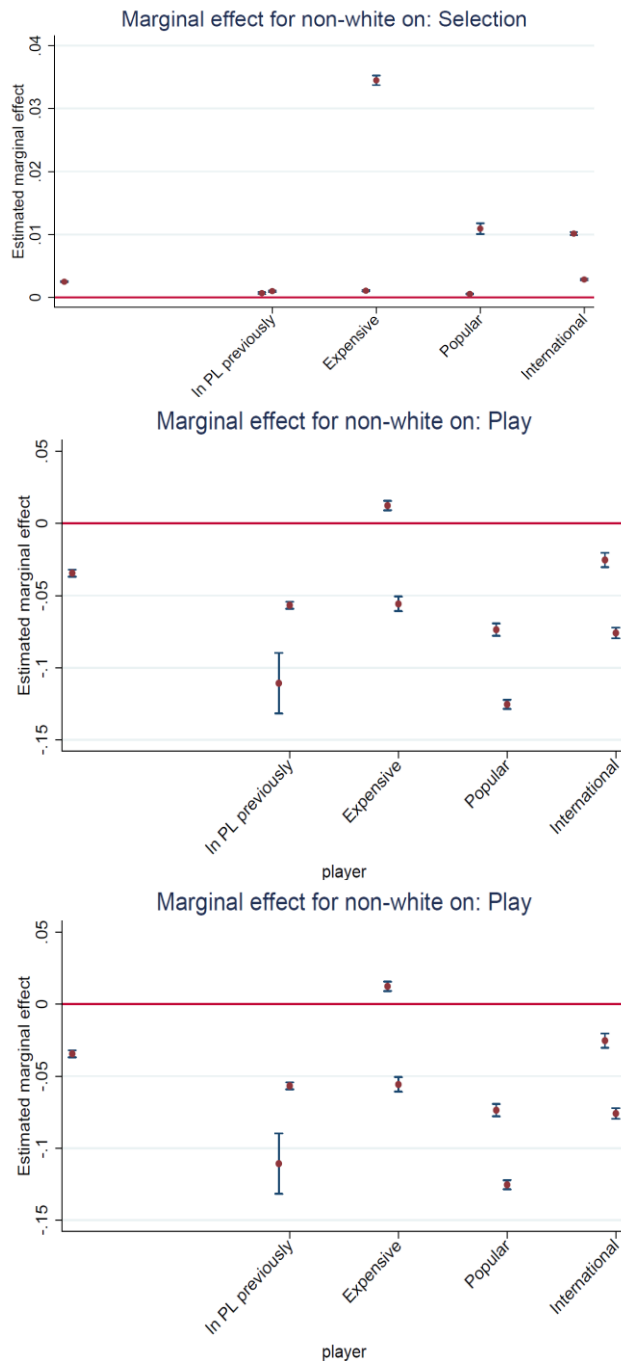
discrimination in the squad selection were ambiguous, for these decisions we find consistent discriminations against non-white.

**Figure 7: Effect of Race on Squad Selection in Week 1**



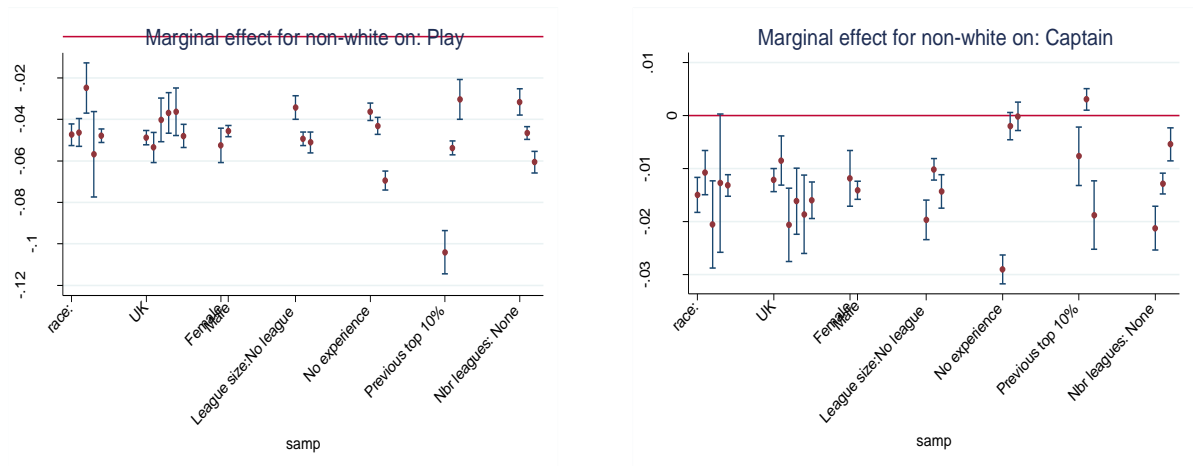
We have seen previously that participants focus their effort on the star performers, in Figure 9 we report the estimates of squad selection, team selection and captancy for star performers in term of price, past performance or simply having played for the national teams. Consistent with the previous evidence, star performers are more likely to be discriminated in team and captancy decisions.

**Figure 9: Racial Discrimination in Week 1 for Star and non-Star Performers**



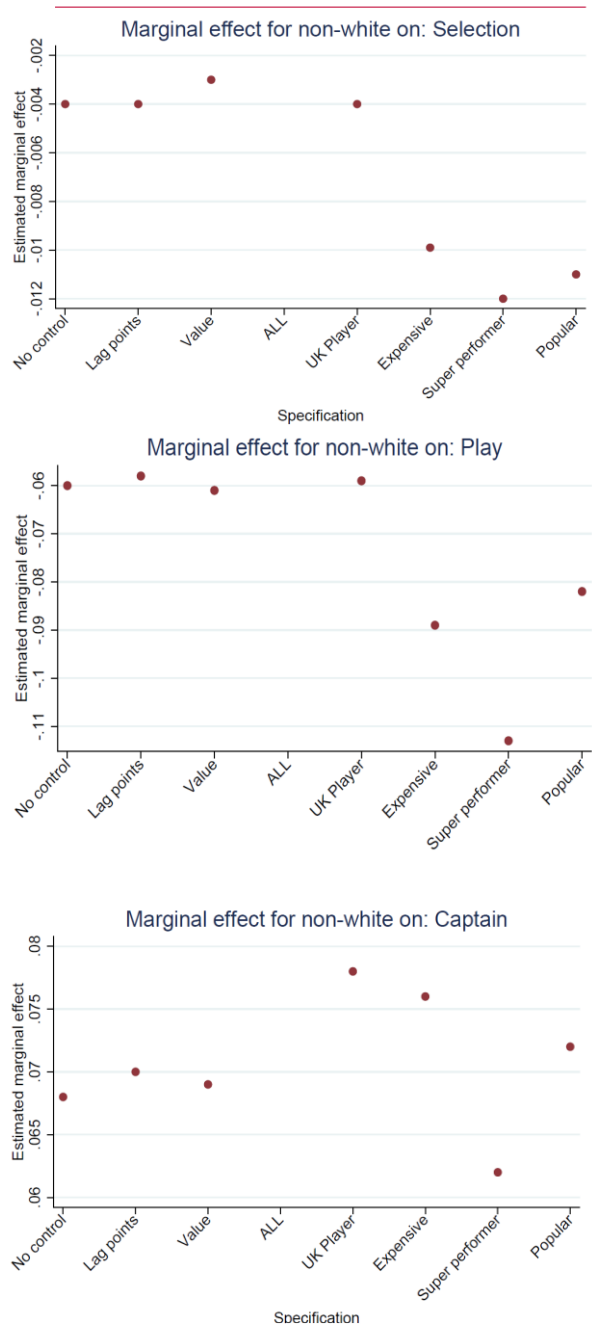
We replicate this analysis about week 1 decisions for the different sub-sample of participants and report it in Figure 10. There is limited variation in the extent of discrimination by participants characteristics.

**Figure 10: Week 1 Discrimination by Participants Characteristics**



Moving on to the rest of the season's decisions, Figure 11 report the estimates on discrimination for different specifications and players. Again we find that stars are the most likely to be discriminated. Since the model includes controls for past selection in the team, this negative coefficient reflects that non-white players are more likely to be dropped out of teams, especially if they are top performers. This is consistent with the evidence that most of the trading is concentrated on the top performer, and suggest that non-white top-performers are more likely to be dropped out of teams than their white counterparts if their performance is not perceived as satisfactory.

**Figure 11: Racial Discrimination in Subsequent Weeks by Players Characteristics**



To be continued...

## 6 Conclusion

We test Becker's prediction that employer racial prejudice affects the employment prospects of marginal workers. Our virtual labour market data with 3 million identical firms

identifies all hiring, firing and promotion decisions over 38 periods. In a context where racial discrimination is possible and driven by taste we find that employers discriminate on racial grounds when deciding who to hire, who to fire, and who to promote. Consistent with Becker's (1957) model of taste-based discrimination, in the context of FPL, non-white footballers need to be more productive than whites to gain a position in the team, to avoid dismissal, and to become team captain.

## References

- Altonji, J. G. and Blank, R. M. (1999) "Race and gender in the labor market", Chapter 48, *Handbook of Labor Economics*, vol. 3, 3143-3259
- Altonji, J.G. and Pierret, C.R. (2001) "Employer learning and statistical discrimination", *Quarterly Journal of Economics*, 464, 1: 313–50
- Al-Ubaydli, O., and List, J. A. (2017). Field experiments in markets. In *Handbook of Economic Field Experiments* (Vol. 1, pp. 271-307). North-Holland.
- Arrow, K. J. (1972) "Models of job discrimination" in A. H. Pascal (ed.) *Discrimination in Economic Life*, pp. 83-102, Lexington, MA: Lexington Books
- Arrow, K. J. (1973) "The Theory of Discrimination" in O. Ashenfelter and A. Rees (eds.) *Discrimination in Labor Markets*, pp. 3-33, Princeton, NJ, Princeton University Press
- Baerg, A. (2009) "Just a Fantasy? Exploring Fantasy Sports". *Electronic Journal of Communication*, 19, 3.
- Baert, S. (2018) "Hiring discrimination: An overview of (almost) all correspondence experiments since 2005." In *Audit studies: Behind the scenes with theory, method, and nuance* (pp. 63-77). Springer, Cham.
- Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016) "Attention discrimination: Theory and field experiments with monitoring information acquisition". *American Economic Review*, 106 (6), 1437-75.
- Becker, G. S. (1957) *The Economics of Discrimination*, Chicago: Chicago University Press
- Bendick, M., Jr. (2007) "Situation Testing for Employment Discrimination in the United States of America", *Horizons Strategiques*, 3(5): 17-39

- Bertrand, M. and Mullainathan, S. (2004) "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination", *American Economic Review*, 94(4): 991-1013
- Bohren, J. A., Haggag, K., Imas, A. and Pope, D. G. (2019) "Inaccurate Statistical Discrimination", *NBER Working Paper No. 25935*
- Bryson, A. and Chevalier, A. (2015) "Is there a Taste for Racial Discrimination Among Employers?", *Labour Economics*, 34: 51-63
- Cahuc, P., Carcillo, S., Minea, A. and Valfort, M.A. (2019) "When Correspondence Studies Fail to Detect Hiring Discrimination", IZA Discussion Paper No.
- Charles, K. K. and Guryan, J. (2008) "Prejudice and Wages: An Empirical Assessment of Becker's *The Economics of Discrimination*", *Journal of Political Economy*, 116(5): 773-809.
- Charness, G., & Kuhn, P. (2011). Lab labor: What can labor economists learn from the lab?. In *Handbook of labor economics* (Vol. 4, pp. 229-330). Elsevier.
- Di Stasio, V. and Heath, A. (2019) "Racial Discrimination in the British Labour Market", *University of Oxford, Mimeo*
- Dovidio J. F, and Gaertner, S.L (2000) "Aversive racism and selection decisions", *Psychological Science*, 11(4):315–19
- Goff, B. L., McCormick, R. E., and Tollison, R. D. (2002) "Racial Integration as an Innovation: Empirical Evidence from Sports Leagues", *The American Economic Review*, 92, 1: 16-26
- Guryan, J. and Charles, K. K. (2013) "Taste-based or Statistical Discrimination: The Economics of Discrimination Returns to Its Roots", *The Economic Journal*, 123, F417-F432



- Hart, A., P. Whalen, L. Shin, S. McInerney, H. Fischer and Rauch S. (2000) Differential response in the human amygdala to racial outgroup vs ingroup face stimuli. *Neuroreport*, 11, 2351-2354
- Hedegaard, M. S. and J.-R. Tyran (2018) "The Price of Prejudice," *American Economic Journal: Applied Economics*, 10: 40-63.
- Kahn, L. M. (2009) "The Economics of Discrimination: Evidence from Baseball", *IZA Discussion Paper No. 3987*
- Lang, K., and Lehmann, J. Y. K. (2012). "Racial discrimination in the labor market: Theory and empirics". *Journal of Economic Literature*, 50(4), 959-1006.
- Laouénan, Morgane, and Roland Rathelot. (2022) "Can information reduce ethnic discrimination? Evidence from Airbnb." *American Economic Journal: Applied Economics* 14, no. 1: 107-32.
- Levitt, S (2004) "Testing Theories of Discrimination: Evidence from *The Weakest Link*." *Journal of Law and Economics*, 47(2): 431–52
- List, J. A. (2004) "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field", *The Quarterly Journal of Economics*, 119, 1: 49-89
- Neumark, D. (2018) "Experimental Research on Labor Market Discrimination", *Journal of Economic Literature*, 56, 3, 799-866.
- Parsons, C., Sulaeman, J., Yates, M. and Hamermesh, D. (2011) Strike Three: Discrimination, Incentives and Evaluation. *American Economic Review*, 101, 1410-1435
- Phelps, E., K. O'Connor, W. Cunningham, E. Funayama, C. Gatenby, J. Gore and Banaji, M. (2000) Performance on Indirect Measures of Race Evaluation Predict Amygdala Activation. *Journal of Cognitive Neuroscience*, 12, 729-738

- Phelps, E. (1972) "The Statistical Theory of Racism and Sexism", *American Economic Review*, 62, 4:659-661
- Pope, D., and Sydnor, J. (2011) "What's in a picture? Evidence of Discrimination from Prosper.com", *Journal of Human Resources*, 46, 1: 53-92
- Rosen, S. (1981) "The economics of superstars", *American Economic Review*. 71; 845–858
- Rosen, S. and Sanderson, A. (2001) "Labour Markets in Professional Sports", *The Economic Journal*, 111, F47-F68
- Weichselbaumer, D. (2016) "Discrimination against female migrants wearing headscarves", *IZA Discussion Paper No. 10217*

# Appendix

To be completed