

DISCUSSION PAPER SERIES

IZA DP No. 15473

**Discrimination in a Rank Order Contest.
Evidence from the NFL Draft**

Ian Gregory-Smith
Alex Bryson
Rafeal Gomez

AUGUST 2022

DISCUSSION PAPER SERIES

IZA DP No. 15473

Discrimination in a Rank Order Contest. Evidence from the NFL Draft

Ian Gregory-Smith

University of Sheffield

Alex Bryson

University College London and IZA

Rafeal Gomez

University of Toronto

AUGUST 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Discrimination in a Rank Order Contest. Evidence from the NFL Draft

This paper examines discrimination in the NFL draft. The NFL is a favourable empirical setting to examine the role of skin colour because franchise selectors are required to make rank-order judgements of players based on noisy signals of future productivity. Since wages are tightly related to the rank-order of the draft for the first four years of a player's career, even if discrimination plays only a marginal role in selection, there could be a large discriminatory impact. We observe large unadjusted racial differences in drafting. However, much of the variation is explained by Black and White players selecting into different playing positions. Conditional upon a large set of control variables, including athletic performance at a marquee selection event (the NFL combine), we do not find robust evidence of racial discrimination in NFL drafting between 2000 and 2018.

JEL Classification: J15, J24

Keywords: discrimination, race, NFL

Corresponding author:

Alex Bryson
Social Research Institute
University College London
55-59 Gordon Square
London WC1H 0NU
United Kingdom
E-mail: a.bryson@ucl.ac.uk

1 Introduction

Entry into professional careers usually requires potential employees being invited to a selection event such as an interview or an assessment day. Performance at this event, combined with other information about the candidate then determines the rank order of candidates and who is offered the job. In some professions such as a law, medicine and academia, graduates will try out for several potential employers and their rank order will be strongly correlated with their starting salary. By this process, top candidates are sorted into the highest paying jobs.

This process is at play in the selection of professional American football players out of college via the NFL draft. College players are invited to a selection event, called the NFL Combine, at which their performance over a series of tasks is observed. Then at the NFL draft, the 32 NFL franchises take turns in picking the candidates with their wages following the order of the pick ¹. After seven rounds of picking players, the remaining college player pool is considered undrafted.

Ideally, gender, age or skin colour² should play no role in the hiring of candidates, but there is evidence from the economics literature to suggest that this isn't the case (Bertrand and Mullainathan, 2004; Goldin and Rouse, 2000; Neumark, 2020). Either because of prejudice (Becker, 1957), or 'statistical' processes that break along ageist, racial or gendered lines (Arrow, 1972; Phelps, 1972) it is possible that the non-productivity related characteristics of the candidate impacts hiring. We argue that identifying racial bias in the hiring process is important not just because we want hiring processes to be fair. If there is bias in the hiring process, a regression of the observed wage gaps between workers of different races may underestimate the full extent of discrimination in the labour market. However, outside of professional sports, identifying racial bias in hiring is difficult (Parsons et al., 2011). It is usually not possible for researchers to observe an accurate signal of a candidate's potential productivity. As such, it is usually not possible to differentiate between competing interpretations for observed differences between groups. Additionally, acquiring information on unsuccessful candidates is necessary to measure the success rates accurately (Neumark and Rich, 2019). In our setting, information on successful and unsuccessful candidates is publicly broadcast.

This paper merges information from a number of publicly available sources to create a unique data set that tracks college players through the NFL Combine and NFL Draft and into their first contract as a professional football player for an NFL franchise. By observing publicly available images of players, we manually assign players to one of three categories: 'Black', 'White' or 'Non Black Non White'. We observe large unconditional differences in drafting based on these groupings. However, much of the variation can be explained by Black and White players selecting into different playing positions. Additionally, after controlling for a full set of productivity measures we do not find robust evidence of racial discrimination against Black players in drafting. In fact, our results suggest there is an unexplained advantage for Black players across most playing positions, other than at quarterback. This is the case both at the extensive margin of drafting vs not drafting and also on the intensive margin of how early in the selection process a player is drafted i.e. their rank position within the draft.

¹Please refer to the appendix for details

²We use the term 'skin colour' in this paper to describe differences in NFL rankings and draft position between Black and White players in the NFL. This journal, and the economics literature at large, uses the term 'race' and the variables we define in this paper can be thought of as racial categories. But note that 'race' is to be understood here as a social construction (there are no such things as 'biological races' (Sussman, 2014)). Put simply, we observe a player's skin colour and assign them to one of three categories: 'White', 'Black' or 'Non White Non Black'.

There are some aspects of the data that could be a concern for Black participants. The first is that Black quarterbacks are overlooked in favour of White quarterbacks with the same physical characteristics, albeit this difference is small and only marginally statistically significant. It is only when the measure of cognitive ability, known as the Wonderlic test³, that Black quarterbacks achieve parity. In other words, it is because Black quarterbacks under perform on this test relative to White quarterbacks, that White quarterbacks have greater success on average in the draft. The average performances on this test mirror those found in the broader US population. Nevertheless, the question remains whether this is a unbiased test of quarterback playing ability (Lyons, Hoffman and Michel, 2009) and indeed the NFL will no longer administer this test from 2022 onwards (See <https://www.nytimes.com/2022/03/02/sports/football/nfl-wonderlic-test.html>).

Another concern is that players who are neither Black or White, typically players with Hispanic, Pacific Islands, or Asian backgrounds under perform at the quarterback position given their observed characteristics. We should stress that the statistical power of this result is low because there are only eleven Non Black Non White quarterbacks in the 18 years of our sample. We can't say if this result would hold had their been more observations. However, the very fact that there are so few quarterbacks from these groups in the sample is perhaps an indication that such players face barriers to entering the NFL.

2 Background

Economists usually begin from the position that hiring is a hidden type problem. Ex ante, candidates have private information over their own future productivity that is hidden from employers. Even after hiring, since observed productivity typically contains some noise, it may take several time periods before employers are confidently able to identify and remove low productivity workers. This would be costly for the employer. Therefore, it is profitable for employers if they can obtain information to improve their ranking of candidates. Interviews and assessment days therefore represent relatively low cost mechanisms to obtain a signal of future productivity. In our context, the signal is obtained through the physical tests at the NFL Combine along with playing statistics in college. This signal can be combined with other information about the candidate to arrive at a rank order that maximises future productivity.

However, the other information may include demographic averages, as it does in models of 'statistical' discrimination (Arrow, 1972; Phelps, 1972). In such models, it is assumed that the employer knows the distribution of productivity for each group (by age, gender, race etc) and they can use this information to form a more accurate prediction of expected productivity than would be obtained if they used the individual level productivity signal alone. In short, if the observed average level productivity of one demographic group is higher than the other, candidates from that group will be ranked higher given equal individual productivity signals⁴.

A recent theoretical development in a similar spirit to statistical discrimination are models of biased beliefs in hiring (Bohren, Imas and Rosenberg, 2019; Schwartzstein, 2014). The key dif-

³Potential NFL players since 1970 have had administered an intelligence measurement called the Wonderlic Personnel Test at the NFL Scouting Combine. The test is used to measure players' aptitude for learning and problem solving. The possible score range is 1 to 50. The average football player scores around 20 points and scores vary by position. See: <https://fivethirtyeight.com/features/how-a-multiple-choice-test-became-a-fixture-of-the-nfl-draft/>

⁴Other moments of the productivity distribution might also be relevant. For example, a risk averse employer may prefer candidates from a group with lower variance, all else equal

ference to the traditional statistical discrimination model is that the employer's understanding of the productivity distribution by race may be incorrect. Due to stereotyping or historical performances, a franchise may believe, for example, that White players outperform Black players at quarterback and make hiring decisions accordingly. However, this assessment of the productivity distribution may be outdated and should be subject to revision as new information on player performances emerges. Franchise selectors who hold on to outdated stereotypes may fail to update their assessment of the productivity distribution, even when new statistical information becomes available.

There is qualitative evidence consistent with stereotyping. In Dufur and Feinberg (2009) the authors conducted interviews with players invited to the NFL combine. While they found no evidence of overt discrimination, players from ethnic minorities reported practices that were consistent with racial stereotyping. For example, Black and Polynesian players were questioned about family relationships and off-field (potentially criminal) behaviour to a greater extent than White players. In Mercurio and Filak (2010) a text analysis was conducted on the media commentary of Black and White college quarterbacks. A large difference was notable in how Black and White quarterbacks were described. Mental strength and decision making were associated with White quarterbacks while Black quarterbacks were more frequently described in terms of their physical attributes.

Alternatively, (or in addition) to statistical discrimination, other forms of discrimination could be at play. The most egregious mechanism is that of simple employer prejudice known as 'taste-based' discrimination (Becker, 1957). In these models, employers are willing to pay for their prejudice in the form of higher wages for White employees. If White and Black employee wages are equalised (for example due to a collective bargaining agreement), then fewer Black employees will be hired. With wages tightly following the draft order, we would expect taste discrimination to be seen in the rank order of the draft, or the selection of marginal White draftees over Black draftees. This was certainly the case in the past, given the known cases of historical racism in American football. For example, from 1869 to 1945, Black players were effectively excluded from playing, with only a handful of exceptions (Levy, 2003; Smith, 1988).

In a perfectly competitive market, such discriminating employers can not survive the forces of competition, but in the NFL, there are significant restrictions on competition that may allow prejudiced employers to exercise their taste for discrimination (Neumark, 1988). In the name of 'competitive balancing', there is revenue sharing between franchises, restrictions on 'rookie' contracts, a hard salary cap, and a reverse order of finish college draft. These competitive balancing mechanisms have been the source of much academic discussion (see Késenne (2014); Szymanski (2006)) but a further unintended consequence of them could be that they permit employer discrimination.

A subtle variation on taste-based discrimination locates the source of the prejudice with the employer's customer base, in this case football fans. Put simply, a majority White fan base may be willing to pay more to see a White quarterback, as the 'face of their franchise'. A profit maximising NFL franchise might be expected to reflect their customers' preferences in the NFL draft, even if those preferences are prejudiced. There is evidence from Kahn (1992) that White and Black player wages vary according to the racial demographic of the metropolitan area from where the fans are drawn. A related mechanism is that fans may reduce the productivity of players from minority groups through racist intimidation. A recent paper by Caselli, Falco and Mattera (2022) uses the absence of fans in association football due to Covid restrictions to estimate the reduction in productivity of players with African heritage caused by racist fans. Again, a profit maximising franchise would be expected to factor in the lower productivity arising from fan abuse, even if the franchise selectors themselves are not prejudiced.

A final concern is that players may face discrimination from other players, known as ‘employee discrimination’. Given the potential for productivity spillovers between players (e.g. a good receiver will raise the productivity of a quarterback), the possibility of employee discrimination lowers the expected productivity of a player from a minority group. However, we argue that in our context, while not impossible, this is unlikely to be widespread because by the time players enter the NFL they have been playing in racially diverse teams for several years. Additionally, to our knowledge, we don’t have any anecdotal evidence of this having occurred during our sample period.

2.1 Empirical literature

In testing for discrimination in the NFL, the literature has, by and large, taken the traditional regression based approach of Mincer (1970). Essentially, a wage equation is specified with a dummy variable identifying race and other variables controlling for other key determinants of wages such as human capital and prior productivity⁵. When applying this approach to NFL data, it becomes apparent that the race variable is typically insignificantly different from zero. That is to say, at the mean, there is no unexplained wage gap between White and Black NFL players (Berri and Simmons, 2009; Burnett and Scyoc, 2015; Ducking, Groothuis and Hill, 2014).

Nevertheless, the race identifier can be interacted with performance statistics to reveal differences in how productivity between the races is rewarded in the labour market. In particular, (Berri and Simmons, 2009) find that rushing yards by Black quarterbacks are not compensated in their sample. The authors also find some evidence of wage gaps at points of the distribution away from the mean, particularly at the top end of the wage distribution. This result is also seen in (Keefer, 2013), who identifies a wage gap for linebackers at all points of the distribution but particularly at the top and bottom deciles. However, the effect does appear to be confined to certain positions. Across a range of different NFL positions, Ducking, Groothuis and Hill (2014) do not find wage gaps when adopting a quantile approach.

Our test of discrimination takes a different approach to the literature above. Rather than examine the observed wage distribution, we want to focus on the rank ordering of candidates in the NFL draft, those players having come through college football and the NFL Combine selection event. Since player wages during the first four years of the playing career, (their ‘rookie’ contract) are very tightly associated with their draft position, we argue that the relevant point of assessment is the draft. Further, a significant number of NFL players do not survive in the league beyond their rookie contract (Volz, 2017).

While we bring the largest sample in the literature to date, our paper is not the first to study the NFL draft. Hendricks, DeBrock and Koenker (2003) examine statistical discrimination in the draft with respect to players from less visible college programs (non Division 1A). If these players are drafted early, they tend to have better careers than equivalent players from high profile schools. This implies that players from low profile schools have to be better than their equivalents to draft early, reflecting statistical discrimination. However, in later rounds the effect is reversed with Division 1A players ranking lower than equivalent non Division 1A players. The authors interpret this as support for an ‘option value’ effect. Because a franchise can eliminate poor performers ex

⁵As noted by Simmons (2021) in a recent review, the literature has generally not applied the Oaxaca–Blinder decomposition because of the insufficient number of observations of players from the minority group to estimate two separate equations.

post, it becomes profitable for franchises to take a risk on a less visible player in the hope that they uncover a hidden ‘star’.

Massey and Thaler (2013) present evidence that franchises overvalue top picks in the draft. Evidence for this claim is based on the difference between the compensation cost of the pick relative to what an equivalent player would cost in free agency. As noted by Berri and Simmons (2011) this approach requires the strong assumption that the draft itself does not impact future pay. Nevertheless, the broader claim that the draft is an imprecise process in terms of correctly identifying future talent is consistent with Berri and Simmons (2011), particularly at the quarterback position. While drafting early is predictive of playing time, conditional upon playing time, early picks do not always outperform later ones. Further, the variables that are predictive of where a player will draft do not translate to playing success. Wolfson, Addona and Schmicker (2011) revisit this question and although they agree that the process is imprecise and that variables predicting draft and combine success are not strongly predictive of playing performance, they offer a different interpretation to Berri and Simmons (2011). The authors argue that franchises aggregate all the pre-draft information about a player, including non-quantitative information such as scouting reports. This results in the draft position containing significant information about playing suitability, which is then reflected in more game time for the players.

Berri and Simmons (2011) also do not find significant differences in draft rank order between Black and White quarterbacks. Gill and Brajer (2012) continue this investigation in more detail by examining quarterbacks in the draft and focusing on the measure of cognitive performance at the Combine, known as the Wonderlic test. If franchises were treating the Wonderlic signal of cognitive ability as a stronger signal for White quarterbacks than for Black quarterbacks, this could be interpreted as evidence of statistical discrimination. However, the authors find that this is not the case but rather the marginal effect of the Wonderlic for Black and White quarterbacks is equal. In contrast, Conlin and Emerson (2005) do find evidence of discrimination in the draft by separating the drafting decision from the playing decision. While drafting early greatly improves the chances of playing in the NFL, there is no guarantee and players can be dropped before the start of the season. Conditional upon drafting position (and other variables), non-White players are shown to be more likely to play and start games than White players. The authors interpret this result as evidence of discrimination in the draft against non-White players by assuming the decision to play is race neutral. If this assumption holds, coaches who pick the playing team are correcting the drafting bias from the front office, who should have been drafting the playing non-White players earlier in the draft.

Our contribution to this literature is to provide a more complete picture of the role of race in the NFL Draft. We bring together information from several sources to form the largest dataset to date with information on both successful and unsuccessful applicants. In doing so, we emphasise the importance of playing position and the sorting of players by race into different playing positions.

3 Data

Our dataset is an individual level dataset comprising the full population of Combine participants from the year 2000 to 2018. We create three player categories, Black (N=4,206), White (N=1,669) and Non Black % Non White (N=361) by observing the images of each player which are in the

public domain⁶.

As shown in Table 1, there are 6,770 participants of which 3,895 (57.5%) are successful in the draft (were drafted in one of the seven rounds). Success rates vary between position reflecting the variation in the supply and demand for positions. 54% of quarterbacks draft, while as many as 63% of defensive linemen draft and as few as 26% of special teams (Kickers and Punters) draft. The majority of draftees are Black but there is large variation between positions. There are more White draftees than Black draftees at the position of quarterback, offensive linemen (which includes the tight end position) and special teams. In all other positions there are more Black draftees.

Table 1: Drafting by Position

	All	QB	Oline	Dline	Dcover	LB	RB	WR	Sp
N	6,770	379	1,469	1,064	1,202	789	709	923	235
N Drafted	3,895	203	834	672	742	489	383	512	60
<i>Black</i>									
N	4,206	90	554	740	1,006	519	565	730	2
N Drafted	2,770	44	367	516	684	365	334	459	1
<i>White</i>									
N	1,669	278	682	150	55	147	60	85	212
N Drafted	962	157	419	109	39	105	35	41	57
<i>Non Black Non White</i>									
N	361	11	109	85	38	36	33	39	10
N Drafted	163	2	48	47	19	19	14	12	2
<i>% of drafted are:</i>									
Black	71.12%	21.67%	44.00%	76.79%	92.18%	74.64%	87.21%	89.65%	1.67%
White	24.70%	77.34%	50.24%	16.22%	5.26%	21.47%	9.14%	8.01%	95.00%
<i>Difference</i>	46.4pp	-55.7pp	-6.2pp	60.6pp	86.9pp	53.2pp	78.1pp	81.6pp	-93.3pp
SE	0.01	0.08	0.03	0.01	0.00	0.01	0.01	0.01	0.18

1. Sample is NFL Draft Combine participants from 2000-2018

2. *Difference* is the percentage point difference in the proportion of drafted players who are Black vs the proportion of drafted players who are White. SE is the standard error of this proportion. All differences are statistically significant at conventional levels.

However, we can not infer from these numbers alone whether or not Black and White players are more favoured in the draft at different positions. This is because the number of Black and White participation also varies dramatically across positions. In all the positions where there are more successful Black (White) draftees, there are more unsuccessful Black (White) participants. For example, although $57/60 = 95\%$ successfully drafted Special Team players are White, the vast majority of unsuccessful Special Team players are also White ($155/156 = 94\%$). Whatever is causing the vast majority of Kickers and Punters to be White, it is happening prior to participation in the Combine⁷. Conditional upon participation in the Combine, White and Black players are

⁶Classification of players into the three categories was primarily conducted by research assistants Michael Muir & Justin Zelnicker, with oversight from the authors. A random sample from each data inquirer was analysed and no statistically significant differences in classification frequency were present. There were 11 players who we could not confidently classify into either of the three categories and these were dropped from the analysis.

⁷This is not to say the draft, or decisions after the draft in the NFL are necessarily colour blind at this position. It is quite possible that anticipation of unequal treatment at this position gives rise to unequal selection into this position. What is being claimed here is that this process, or whatever else drives the unequal selection is taking effect prior to the

equally likely to be successful at this position.

Nevertheless, the possibility of unequal drafting success rates remain at other positions. This is the focus of Table 2. Overall, of those drafted, the difference in Black-White percentage points is 46.4pp higher for Black players. The equivalent statistic in the undrafted pool is 25.3 percentage points higher for Black players. Therefore, there is a statistically significant difference in these differences of 21.1pp in favour of Black players. Altogether, Black players are more likely to draft, even controlling for the greater number of Black participants at the Combine.

Turning attention to specific positions (reading along the DiD row), Black players are favoured at defensive line, defensive cover, linebacker, running back and wide receiver. The only position where White players are favoured is the position of quarterback, to the tune of 13.05 percentage points. However, this number falls short of statistical significance because of the low number of Black participants at the quarterback position.

Table 2: Black/White Differences in Drafting Success Rates

	All	QB	Oline	Dline	Dcover	LB	RB	WR	Sp
<i>Drafted B vs W</i>									
Difference	46.4pp	-55.7pp	-6.2pp	60.6pp	86.9pp	53.2pp	78.1pp	81.6pp	-93.3pp
SE	0.01	0.08	0.03	0.01	0.00	0.01	0.01	0.01	0.18
<i>Undrafted B vs W</i>									
Difference	25.36pp	-42.61pp	-11.97pp	46.68pp	66.52pp	37.33pp	63.19pp	55.23pp	-88.00pp
SE	0.01	0.08	0.04	0.02	0.00	0.02	0.01	0.01	0.11
DiD	21.06	-13.05	5.73	13.88	20.41	15.84	14.88	26.41	5.33pp
SE	0.01***	0.11	0.04	0.02***	0.01***	0.03***	0.01***	0.01***	0.20

1. Sample is NFL Draft combine participants from 2000-2018.
2. DiD denotes the difference in the differences of drafting success rates between Black and White participants. DiD controls for the unequal numbers of Black and White participants at each position. A positive (negative) number indicates Black participants face a higher (lower) drafting success ratio on average than White participants, irrespective of the unequal number of Black and White participants at each position.
3. Statistical significance is shown at (*10%, **5%, ***1%)

The differences outlined in Table 2 are unconditional differences in drafting success rates between Black and White players. In common parlance, these can be thought of as race *gaps* in drafting success, in the same sense that the gender wage gap is the unconditional mean difference in wages between two groups, men and women. As with wage gaps, our race gap in drafting success by position does not imply discrimination in the draft. This is because after conditioning on other variables that determine the drafting success, it may be that there are no outstanding gaps to explain. Equally, the absence of a race gap does not imply the absence of racial discrimination; gaps may emerge after controlling for other determinants of drafting success.

Yet, that is not to say the presence of a race gap in drafting is uninformative on the matter of discrimination in the draft. Rather it provides the context to understanding the avenues through which discrimination may be occurring. In our case, the presence of a positive race gap for Black players overall and at the specific positions of Offensive Line, Defensive Cover, Running Back and Wide Receiver narrows the range of possible inferences. First, prejudice becomes less likely because pure malice should be independent of the position a player occupies. If a taste-based

draft.

argument is to be maintained, it is more likely to be the stereotyping variant, whereby franchises fail to correctly update their beliefs of a player’s expected productivity, given their position and race.

Second, to hold to the position of no discrimination, it must be that Black participants in the Combine at these positions are drawn from a better distribution of potential Combine invitees than White participants. This might be because the distribution of Black players is simply better than White players in the population at these positions, in which case the no discrimination claim may be plausible. Or Black college players at these positions may have faced a tougher selection process to be invited to the Combine, which could imply racial discrimination prior to the Draft. Black players on average may be less well known and / or attend high schools that do not promote as aggressively to Combine selectors. Similarly, the unconditional advantage White quarterbacks have in the draft could be justified if it can be shown that observable characteristics that make a successful quarterback are more frequently observed among White combine participants. Alternatively, it could go the other way. It may be that after observing predictors of quarterback success, Black players are even more disadvantaged than implied by the unconditional differences. What is required is a model of the process by which players are drafted in order to understand why drafting success gaps may be observed in the data. We now turn our attention to this matter.

3.1 Empirical Strategy

3.1.1 Probit

The starting point for our empirical specification is to model the extensive margin of being drafted or not drafted. We begin with a binary Probit:

$$p_i = Pr[y_i = 1|X] = \Phi(X_i'\beta) \tag{1}$$

where p_i is the probability that NFL Combine participant i is drafted conditional upon a vector X of explanatory variables, Φ is the usual cumulative function for the standard normal distribution and β is the vector of estimated coefficients. In all tables, we report average marginal effects.

Our two main explanatory variables of interest are the indicator variables ‘Black’ and ‘Non White, Non Black’ which identify the marginal effect of race on drafting success relative to White players. Central to the quality of our estimating equation (1) is whether or not the vector X captures the draft selection process. To this end, we split our analysis across three key signals of future playing success. First, the ‘physical signal’ given by the recorded performances on the physical tests at the NFL Combine. Second, with respect to the quarterback position only, the cognitive signal given by performance on the Wonderlic test. Third, the signal of future productivity, generated by playing statistics during their college playing career years.

We also control for other variables that may be reasonably determining whether or not a Combine participant is successful in the draft. In particular, we have a set of year fixed effects and position fixed effects. We argue controlling for position is particularly important since the descriptive evidence from tables 1 and 2 could suggest that being Black is an advantage overall, beneficial in some positions but neutral or detrimental in others. Note that by controlling for position, we are limiting the potential mechanisms of racial discrimination to operate only in the draft event, with player position already given. The investigation of discriminatory mechanisms that occur

prior to the draft and impact player selection into those positions is beyond the scope of this study. Creating the final dataset for the analysis required merging data from several sources.

3.1.2 Tobit

While the extensive margin of drafting could be considered to be the fundamental discontinuity that will shape the likelihood of a successful career in the NFL, there are also substantial differences between drafting high and low in the draft. The number 1 draft pick in 2021 signed a contract worth a total of \$37.7M, whereas a draftee in round seven signed a contract with a total value of approximately \$3.5M. Therefore, it is important to consider the determinants of the intensive margin in terms of how high a player drafts.

The drafting position of a player is an ordinal rank, running from first to last, after seven rounds of 32 franchise picks⁸. For the purposes of capturing the intensive margin of this process, we approximate the ordinal rank of the draft with a continuous variable of drafting position. Because this variable is censored, we use a Tobit:

$$y_i^* = (X_i' \beta) + \varepsilon_i \quad (2)$$

$$y_i = \begin{cases} y_i^* & \text{if } 1 \leq y_i^* \leq 261 \\ 1 & \text{if } y_i^* \leq 1 \\ 261 & \text{if } y_i^* \geq 261 \end{cases}$$

The variable y_i^* is the unobserved latent variable and we observe y_i as the draft position. This variable is that is censored from above at 1, because the best a player can draft is in first place, even if this player was many times better than the rest of the field. The variable is also censored from below, because if a player does not draft, we only observe that they are ranked one position worse than the final pick in the draft.

A variant of the Tobit described above is to consider the within position ranking of players. As shown in Table 1, there is variation between positions in drafting success. For example, the quarterback position is the most critical position on the field, contributing the most to the outcome of a game. Highly rated quarterback prospects are likely to outrank highly rated players from other positions. At the other extreme, there have only been three punters who have drafted in the 1st round, the highest being 17th position. Additionally, when making rank choices, franchises do not only consider the quality of the prospect but gaps in their squad that need filling. Therefore, franchises will create their own internal ranking of draft prospects within position.

If candidates with a particular skin colour are being favoured in the draft, we can expect this to be most acutely felt within position. To allow for this, we can adjust the dependent variable according to the within position rank of each player in each year of the draft. This also allows the explanatory variables measuring the productivity signal received by the franchise to be more accurately tailored to the particular position of interest. For example, measures of passing performance (passing yards, completions, passing touchdowns, interceptions etc) are only relevant for players who have thrown in college.

⁸There is variation between years in the number of franchises and therefore the number of picks

3.1.3 Identification

We are adopting a regression based approach, with the intention that we adequately capture the process by which players are drafted. Regression based approaches are notoriously vulnerable to identification issues. In this case, our main variable of interest is a player's skin colour (which is not a choice), so we are free of many of the usual problems associated with an endogenous explanatory variable.

Nevertheless, we should be clear on the following. First, if there are unobserved variables correlated with skin colour that also impact drafting, these will impact the reported marginal effects on race. While we are able to control for a rich set of signals that likely determine drafting success, it is possible that there is information available to franchise selectors but unavailable to us. For example, off-field disciplinary offences may be unequally distributed between Black and White participants. To the extent that such information is correlated with skin colour and impacts drafting success, our estimates will be contaminated by this information.

Second, we are making the assumption that the observed signals of future productivity are themselves independent of racial bias. If Franchises are using racially biased signals, the absence of an effect on the skin colour identifier itself, would not imply a colour blind process. In particular, the Wonderlic test of cognitive ability given to participants at the Combine is controversial, (Gill and Brajer, 2012; Jensen, 1977; Lyons, Hoffman and Michel, 2009) and the NFL has decided no longer to administer the test from 2022 onwards.

Third, in the specifications presented below, we do not allow the signal to vary in its impact by skin colour. For example, it is possible that the marginal impact of a quick 40 second dash time is different for Black and White players. This third point is easily addressed with interaction terms. We do this as a robustness exercise and find no differences of note.

4 Results

Results using the full sample of NFL combine participants from all positions are given in Table 3. Column (1) shows the unconditional gap in drafting between Black and White participants, with Black participants being 21.2 percentage points more likely to draft. This mirrors the DiD estimate in Table 2 above. Columns (2) through (5) add controls for position and year fixed effects, the Physical Signal from the Combine tests and the Productivity Signal from the performances at College. As observable controls are added, the estimated average marginal effect declines but remains statistically significant at 10.8 percentage points. This suggests that the better observable signals (on average) by Black participants explain approximately half of the unconditional gap in drafting likelihood.

The estimates reported in column (1) for 'Non Black Non White' do not indicate any unconditional gap in drafting relative to White participants. Recall, this category is much smaller including only 361 from 6770 participants and comprises players with Hispanic, Pacific Island or Asian backgrounds. A penalty for these players of 5.12 percentage points emerges with the most complete set of controls included. In other words, White participants with similar observable characteristics are less likely to draft overall than Black participants but more likely than players who are neither Black or White.

The control variables behave as expected. The set of year and position fixed effects are jointly statistically significant and impact on the size of the estimated marginal effect (column (2)). Franchise selectors use the Physical signal from the Combine (columns (3) and (4)) as well as the Productivity signal (column (5)) from College football. A little care is required in reading individual estimates from the table. For example, heavier players are favoured in the draft but this holding Height and Speed constant. Together, these variables indicate a more powerful athlete.

Estimates on the intensive margin from Table 4, present a similar qualitative picture to those given on the extensive margin in Table 3. Again, the unconditional gap in drafting in favour of Black participants is large, to the tune of 73 places (about 2 and half rounds higher). This is diminished to 41 places with the full set of observable controls but remains statistically significant and would be financially significant for the player involved. The estimates with respect to Non Black and Non White players are consistent with those of Table 3, albeit fall short of statistical significance. Altogether, Black players are significantly favoured by franchise selectors in the draft for reasons that are only partly explained by observable characteristics.

4.1 Quarterbacks

As shown above, Black candidates from the Combine are favoured in terms of drafting success and drafting position. However, as shown in Tables 1 and 2, there is significant variation by position. We are particularly interested in the quarterback position for three reasons. First, we have strong historical evidence of racism specific to this position as well as anecdotal contemporary evidence suggesting that stereotypes surrounding Black quarterbacks continue to carry weight. Second, the quarterback is the highest paid and most influential position on the team. Highly rated quarterback prospects typically command the highest positions in the draft. With the distribution of wages being tightly correlated with draft position but in a non-linear manner (e.g. see Keefer (2013)) if there is discrimination at the Quarterback position against Black candidates it could offset any average wage advantage Black players in other positions experience. Third, we have additional information on cognitive performance through the Wonderlic test that is only given to quarterback prospects. For these reasons, we focus in on the determinants of draft success for quarterbacks separately.

Columns (1) through (3) of Table 5 report the average marginal effects following a binary probit of whether the player drafted or not, while columns (4) through (6) report the results of a Tobit for drafting position. The advantage that Black players experience in other positions is not present at the quarterback position, but neither are they disadvantaged. A small gap emerges when controlling for the Physical and Productivity signals pertinent to quarterback drafting. However, when additionally controlling for the Wonderlic signal of cognitive ability available to Franchise selectors, there is no unexplained gap for Black quarterbacks, either in drafting or in terms of where a quarterback drafts. This implies, and indeed it is the case, that Black quarterbacks under perform White quarterbacks in the Wonderlic (by approximately 1 standard deviation in our sample⁹). This raises the question whether or not the Wonderlic is racially biased against black participants. Although, the answer to that question is beyond the scope of the study here, we can answer the claim as to whether or not Franchise selectors use the Wonderlic differently for Black and White quarterbacks. Figure 1 presents the distribution from the 1st percentile to the 99th percentile of Wonderlic performance and its impact on drafting success. While higher scores on the Wonderlic increase drafting likelihood, and as stated White quarterbacks on average score higher, there is

⁹This is a consistent finding seen as early as 1977 in the literature Gill and Brajer (2012); Jensen (1977).




Figure 1: Impact of Wonderlic on Quarterback Drafting Probability

Note: The Wonderlic scale runs from zero to 50. A score of 11 represents the 1st percentile and 43 represents the 99th percentile. Overlapping confidence intervals for Black and White quarterbacks suggest that there are no statistically significant differences in how Wonderlic impacts drafting success.

no statistically significant difference in terms of the marginal impact of a higher Wonderlic score¹⁰. Altogether, these results do not suggest that being Black is a disadvantage in the draft at the quarterback position.

A different picture emerges for players who are neither Black nor White. There is a large unconditional gap in drafting success and drafting position. This gap is diminished but remains large and statistically significant after controlling for the Physical, Productivity and cognitive signals known to Franchise selectors. However, it should be noted that there are only eleven Non Black, Non White quarterbacks in our sample. The fact that participation is low reduces the confidence that the results are non-random, while at the same time gives pause for thought as to why the participation of Non Black Non White players is so low. One possibility is that these players understand the low likelihood of drafting success at the quarterback position and select out of the competition, perhaps mirroring the experience of Black players 50 years ago.

4.1.1 Within Position Rank

Our final Table 6 shows the impact of our variables on the within position-year rank of the player. That is, where, relative to the competition in their chosen position and year cohort did the player

¹⁰This seems to cohere with what is quantitatively understood to be the use of the Wonderlic score in practice. As noted in a recent article “Wonderlic gives you an area to investigate”, the late New York Giants general manager George Young told the Philadelphia Daily News in 1997. “If a guy doesn’t have a good score on the test, you don’t say he’s not smart. But you go in and investigate and find out [why he scored low]. You go in and talk to his coach. You find out how he did in school. You find out how he retains. If you think he’s a poor reader and did poorly because it was a verbal test, you give him a non-verbal test.”(Siegal, 2015)

rank. This is arguably a more direct competition over player ability that eliminates differences associated with the value of different playing positions. Columns (1) and (2) return the raw and conditional results respectively for all playing positions, while column (3) returns the estimates for the quarterback position. The qualitative picture from the previous analysis is unchanged by focusing on the within position competition. Black players are preferred overall and experience no disadvantage at the quarterback position, conditional upon observables. In the case of the eleven Non Black Non White quarterbacks, these players are ranking lower than expected, given their observed characteristics.

5 Conclusion

Many field experiments and regression-based studies have considered the traditional dimensions of discrimination in labor markets—such as discrimination based on skin colour. These studies nearly always find evidence of discrimination against minorities. The estimates of discrimination in these studies can be biased, however, if there is differential variation in the unobservable determinants of productivity or in the quality of majority and minority groups. Moreover, there is also a potential for bias that runs the other way, in the form of selection effects that may preferentially rank some candidates by dint of their skin colour more highly. This higher ranking may itself produce higher returns later on in careers irrespective of quality/productivity differences.

In this paper by using observable NFL player draft and performance data, we have a study of discrimination that has sufficient information to correct for these two biases. We have data on how highly workers are ranked prior to hiring and have extensive information on player ability and expected productivity.

Altogether, our results do not find robust evidence of hiring discrimination against Black players in the NFL. However, we note three important caveats. First, the absence of evidence of discrimination at the point of hire does not preclude discrimination at other moments in a player's career, either prior to entry or later on as the player enters free agency. Second, the remarkable selection into playing positions prior to hiring may itself reflect racial stereotyping or the anticipation of discrimination. Third, we *do* find some evidence of hiring discrimination against players who are neither Black or White. Recall, this category comprises players with Hispanic, Pacific Island or Asian backgrounds, approximately 5% of the drafting population. While the low numbers of players from these backgrounds reduces the confidence that the results are robust, it also provides pause for thought as to why the participation of Non Black Non White players is so low. One possibility is that these players understand the low likelihood of drafting success and select out of the competition, perhaps mirroring the experience of Black players 50 years ago.

Nevertheless, the absence of discrimination between Black and White professional sportsmen at the point of hire might not be a surprise. Goff, McCormick and Tollison (2002) show that the productivity differential by race in professional basketball and baseball had effectively disappeared by the 1980s, implying the absence of hiring discrimination. Yet in the labour market more generally, correspondence studies suggest substantial and persistent racial discrimination in call backs to interview (Bertrand and Mullainathan, 2004; Heath and Di Stasio, 2019). One rationalisation for the difference in findings could be that the observed call back discrimination is predominately statistical and that the high quality of individual level signals in the NFL and other professional sports settings is sufficient to render the group level signal of limited value.

A key takeaway from our analysis is that the main descriptive differences in the data are a result of sorting into different playing positions by White and Black players. This is happening prior to being assessed for entry into the NFL, most likely in High School. However, why this happens, what the mechanisms are, the extent to which anticipated discrimination is a factor and whether sorting by position will continue to the same extent into the future are open questions for future research. Our work here only suggests that the policy response should be focused earlier than the point of hire at the NFL Draft.

Table 3: Extensive Margin of Drafting Success: Probit

	(1)	(2)	(3)	(4)	(5)
Black	0.212*** (18.03)	0.166*** (11.73)	0.135*** (9.655)	0.139*** (6.418)	0.108*** (7.710)
Non Black Non White	0.0122 (0.454)	-0.0415 (-1.560)	-0.0326 (-1.255)	-0.0462 (-1.273)	-0.0512** (-2.018)
Physical Signal					
Height			-5.81e-05 (-0.00631)	-0.0114 (-0.784)	0.00314 (0.349)
Weight			0.294*** (16.98)	0.340*** (11.30)	0.309*** (18.17)
Speed			0.300*** (25.42)	0.203*** (9.179)	0.298*** (25.77)
Vertical				0.0366** (2.271)	
Bench				0.0419*** (3.478)	
Broad jump				-0.0119 (-0.614)	
Cone				0.0746*** (3.817)	
Shuttle				0.0854*** (4.482)	
Productivity Signal					
College Rank					-0.000701*** (-5.413)
Major School					0.257*** (10.33)
College Performance					0.178*** (10.38)
Year FE	No	Yes	Yes	Yes	Yes
Position FE	No	Yes	Yes	Yes	Yes
Observations	6,762	6,759	6,566	2,974	6,566

1. The sample is NFL Draft combine participants from 2000-2018. Reported estimates are average marginal effects following a Probit. T-statistics of the average marginal effects in parentheses.

2. The variables capturing the physical signal are standardised so that higher values are associated with better performance. College Rank is the within sample historical ranking of the College (lower numbers associated with a stronger college football team). Major School indicates the participant played in a FBS-equivalent conference, or their college played the majority of their games against other rated schools. College Performance is the player's rank in their position according to the playing statistics at College that are relevant to their position.

Table 4: Intensive Margin of Drafting Success: Tobit

	(1)	(2)	(3)	(4)	(5)
Black	-73.16*** (-17.96)	-62.43*** (-13.34)	-49.51*** (-11.23)	-44.36*** (-7.360)	-40.94*** (-9.503)
Non Black Non White	-11.01 (-1.244)	3.456 (0.392)	0.493 (0.0596)	3.683 (0.353)	5.339 (0.663)
Physical Signal					
Height			-1.442 (-0.514)	6.060 (1.516)	-2.653 (-0.974)
Weight			-114.2*** (-20.87)	-127.5*** (-14.88)	-118.3*** (-22.09)
Speed			-115.2*** (-29.00)	-72.59*** (-11.72)	-113.6*** (-29.47)
Vertical				-10.70** (-2.438)	
Bench				-12.25*** (-3.795)	
Broad Jump				-4.809 (-0.926)	
Cone				-26.50*** (-4.921)	
Shuttle				-21.78*** (-4.186)	
Productivity Signal					
College Rank					0.363*** (8.912)
Major School					-102.4*** (-12.24)
College Performance					-57.52*** (-11.27)
Year FE	No	Yes	Yes	Yes	Yes
Position FE	No	Yes	Yes	Yes	Yes
Observations	6,762	6,759	6,566	2,974	6,566

1. The sample is NFL Draft combine participants from 2000-2018. Reported estimates are marginal effects following a Tobit. The dependent runs from 1 through 261, with 1 indicating the highest ranked draftee. Therefore the negative estimates for the variable Black imply an advantage for Black participants.

Table 5: Quarterback Drafting Success

	Drafted: Yes or No			Draft Position		
	(1)	(2)	(3)	(4)	(5)	(6)
Black	-0.0747 (-1.262)	-0.109* (-1.904)	0.0285 (0.406)	9.222 (0.799)	20.20* (1.914)	11.93 (0.833)
Non Black Non White	-0.419** (-2.452)	-0.287** (-2.215)	-0.326** (-2.170)	58.72** (2.008)	47.26** (1.982)	54.43* (1.788)
Wonderlic			0.00917** (2.354)			-1.691** (-2.009)
Physical Signal						
Height		0.214*** (5.121)	0.206*** (4.517)		-43.66*** (-5.597)	-43.41*** (-4.489)
Weight		0.189* (1.769)	0.171 (1.553)		-57.60*** (-3.110)	-58.57** (-2.570)
Speed		0.193*** (4.443)	0.182*** (3.807)		-44.94*** (-5.497)	-45.25*** (-4.546)
Productivity Signal						
College Rank		-0.000949** (-2.315)	-0.00144*** (-3.243)		0.259*** (3.378)	0.332*** (3.263)
College Performance		0.462*** (6.678)	0.249** (2.476)		-75.24*** (-5.283)	-43.14** (-2.059)
Heisman Finalist		0.123* (1.816)	0.150* (1.880)		-54.81*** (-4.530)	-54.27*** (-3.845)
Year FE	No	Yes	Yes	No	Yes	Yes
Observations	379	373	243	379	373	243

1. The sample is NFL Draft combine quarterbacks from 2000-2018. Columns (1)-(3) report average marginal effects following a Probit. Columns (4)-(6) report marginal effects following a Tobit.

2. Wonderlic is a cognitive ability and problem-solving aptitude test. It is administered at the Combine, with the majority of participants being quarterbacks. A Heisman Finalist is an indicator of strong college prospect voted for by sports journalists and previous Heisman trophy winners.

Table 6: Within Position-year Rank

	All Positions		QB
	(1)	(2)	(3)
Black	-15.79*** (-16.92)	-4.819*** (-8.497)	-0.627 (-0.191)
Non Black Non White	-0.586 (-0.289)	1.213 (1.167)	18.51*** (2.653)
Wonderlic			-0.542*** (-2.807)
Physical Signal			
Height		-0.563 (-1.564)	-10.07*** (-4.537)
Weight		-14.91*** (-21.32)	-8.168 (-1.563)
Speed		-14.24*** (-28.87)	-8.803*** (-3.859)
Productivity Signal			
College Rank		0.0423*** (8.054)	0.0891*** (3.814)
Majorschool		-10.98*** (-11.20)	
College Performance		-7.089*** (-10.16)	-11.60** (-2.414)
Heisman Finalist		-18.09*** (-10.87)	-6.588** (-2.036)
Year FE	No	Yes	Yes
Position FE	No	Yes	N/A
Observations	6,762	6,566	243

1. The sample is NFL Draft combine participants from 2000-2018. Columns (1) and (2) report marginal effects within position-year for all positions. Column (3) reports marginal effects for quarterbacks only. Negative numbers imply a positive impact on drafting success, since the number 1 represents the highest pick within position-year.

References

- Arrow, Kenneth.** 1972. "Models of Job Discrimination." *Racial Discrimination in Economic Life*, ed. A. H. Pascal. D. C. Heath.
- Becker, Gary S.** 1957. *The Economics of Discrimination*. University of Chicago Press.
- Berri, David J., and Rob Simmons.** 2009. "Race and the Evaluation of Signal Callers in the National Football League." *Journal of Sports Economics*, 10(1): 23–43.
- Berri, David J, and Rob Simmons.** 2011. "Catching a draft: On the process of selecting quarterbacks in the National Football League amateur draft." *Journal of Productivity Analysis*, 35(1): 37–49.
- Bertrand, Marianne, and Sendhil Mullainathan.** 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 Aislinn, Alex Imas, and Michael Rosenberg.** 2019. "The dynamics of discrimination: Theory and evidence." *American economic review*, 109(10): 3395–3436.
- Burnett, Nancy J., and Lee James Van Scyoc.** 2015. "Compensation Discrimination for Defensive Players: Applying Quantile Regression to the National Football League Market for Linebackers and Offensive Linemen." *Journal of Sports Economics*, 16(4): 375–389.
- Caselli, Mauro, Paolo Falco, and Gianpiero Mattera.** 2022. "When the Stadium Goes Silent: How Crowds Affect the Performance of Discriminated Groups." *Journal of Labor Economics*, forthcoming.
- Conlin, Michael, and Patrick M. Emerson.** 2005. "Discrimination in Hiring Versus Retention and Promotion: An Empirical Analysis of Within-Firm Treatment of Players in the NFL." *The Journal of Law, Economics, and Organization*, 22(1): 115–136.
- Ducking, Johnny, Peter A. Groothuis, and James Richard Hill.** 2014. "Compensation discrimination in the NFL: an analysis of career earnings." *Applied Economics Letters*, 21(10): 679–682.
- Dufur, Mikaela J, and Seth L Feinberg.** 2009. "Race and the NFL draft: Views from the auction block." *Qualitative Sociology*, 32(1): 53–73.
- Gill, Andrew, and Victor Brajer.** 2012. "Wonderlic, Race, and the NFL Draft." *Journal of Sports Economics*, 13(6): 642–653.
- Goff, Brian L., Robert E. McCormick, and Robert D. Tollison.** 2002. "Racial Integration as an Innovation: Empirical Evidence from Sports Leagues." *American Economic Review*, 92(1): 16–26.
- Goldin, Claudia, and Cecilia Rouse.** 2000. "Orchestrating impartiality: The impact of "blind" auditions on female musicians." *American economic review*, 90(4): 715–741.
- Heath, Anthony F., and Valentina Di Stasio.** 2019. "Racial discrimination in Britain, 1969–2017: a meta-analysis of field experiments on racial discrimination in the British labour market." *The British Journal of Sociology*, 70(5): 1774–1798.
- Hendricks, Wallace, Lawrence DeBrock, and Roger Koenker.** 2003. "Uncertainty, Hiring, and Subsequent Performance: The NFL Draft." *Journal of Labor Economics*, 21(4): 857–886.
- Jensen, Arthur R.** 1977. "An examination of culture bias in the wonderlic personnel test." *Intelligence*, 1(1): 51–64.
- Kahn, Lawrence M.** 1992. "The Effects of Race on Professional Football Players' Compensation." *ILR Review*, 45(2): 295–310.
- Keefer, Quinn Andrew Wesley.** 2013. "Compensation Discrimination for Defensive Players: Applying Quantile Regression to the National Football League Market for Linebackers." *Journal of Sports Economics*, 14(1): 23–44.
- Késenne, Stefan.** 2014. *The Economic Theory of Professional Team Sports: An Analytical Treatment*. Edward Elgar Publishing.

- Levy, Alan H.** 2003. *Tackling Jim Crow, Racial Segregation in Professional Football*. McFarland and Co., Inc.
- Lyons, Brian, Brian Hoffman, and John Michel.** 2009. "Not Much More Than g? An Examination of the Impact of Intelligence on NFL Performance." *Human Performance*, 22: 225–245.
- Massey, Cade, and Richard H Thaler.** 2013. "The loser's curse: Decision making and market efficiency in the National Football League draft." *Management Science*, 59(7): 1479–1495.
- Mercurio, Eugenio, and Vincent F Filak.** 2010. "Roughing the passer: The framing of black and white quarterbacks prior to the NFL draft." *The Howard Journal of Communications*, 21(1): 56–71.
- Mincer, Jacob.** 1970. "The Distribution of Labor Incomes: A Survey With Special Reference to the Human Capital Approach." *Journal of Economic Literature*, 8(1): 1–26.
- Neumark, David.** 1988. "Employers' Discriminatory Behavior and the Estimation of Wage Discrimination." *The Journal of Human Resources*, 23(3): 279–295.
- Neumark, David.** 2020. "Age Discrimination in Hiring: Evidence from Age-Blind vs. Non-Age-Blind Hiring Procedures." *NBER Working Paper 26623*.
- Neumark, David, and Judith Rich.** 2019. "Do Field Experiments on Labor and Housing Markets Overstate Discrimination? A Re-examination of the Evidence." *ILR Review*, 72(1): 223–252.
- Parsons, Christopher A., Johan Sulaeman, Michael C. Yates, and Daniel S. Hamermesh.** 2011. "Strike Three: Discrimination, Incentives, and Evaluation." *American economic review*, 101(4): 1410–35.
- Phelps, Edmund S.** 1972. "The statistical theory of racism and sexism." *The american economic review*, 62(4): 659–661.
- Schwartzstein, Joshua.** 2014. "Selective attention and learning." *Journal of the European Economic Association*, 12(6): 1423–1452.
- Siegal, Alan.** 2015. "How A Multiple-Choice Test Became A Fixture Of The NFL Draft." *FiveThirtyEight*.
- Simmons, Rob.** 2021. "Professional Labor Markets in the Journal of Sports Economics." *Journal of Sports Economics*, forthcoming.
- Smith, Thomas G.** 1988. "Outside the pale: The exclusion of Blacks from the National Football League, 1934-1946." *Journal of Sport History*, 15(3): 255–281.
- Sussman, Robert Wald.** 2014. *The Myth of Race: The Troubling and Persistence of an Unscientific Idea*. Cambridge: Harvard University Press.
- Szymanski, Stefan.** 2006. "Handbook on the economics of sport." , ed. Wladimir Andreff and Stefan Szymanski, Chapter 34 The theory of tournaments, 337–341. Edward Elgar Publishing Ltd.
- Volz, Brian D.** 2017. "Race and quarterback survival in the National Football League." *Journal of Sports Economics*, 18(8): 850–866.
- Wolfson, Julian, Vittorio Addona, and Robert H Schmicker.** 2011. "The quarterback prediction problem: Forecasting the performance of college quarterbacks selected in the NFL draft." *Journal of Quantitative Analysis in Sports*, 7(3).

A Appendix

The NFL Draft: Details

The NFL draft is an annual event whereby NFL Franchises pick from a pool of eligible college players. Each Franchise is awarded one pick for each of the seven rounds. It is a 'reverse order of finish' draft which means the worst performing team in the prior year is awarded the first pick in each round, the second worst the second pick and so on until the Superbowl champions pick last (32nd). Each franchise will have their own internal rank order informed by the players' performances at College and at the NFL Combine event which occurs earlier in the year. Franchises will also use written opinions from their professional scouts. Franchises can select a player from any playing position.

Each pick in the draft is an asset that can be traded. Franchises may swap picks (e.g. their first round pick for two later round picks) or trade their picks for veteran players, who are already established in the NFL. Trades may cross years (a first round pick today, might be traded for two picks in later years). Trading can occur during the off-season from March to November (the season runs September to February). NFL franchises use draft trade charts which guide how much a draft pick is worth.

Draftee wages are largely determined by the league wide Collective Bargaining Agreement (CBA) which has been in place since 2011. The CBA assigns a maximum and minimum value for each pick position in the Draft. The exact formula for draftee wages is not disclosed but negotiations over wages typically conclude quickly after drafting upon which the agreed wage package is observed. Top draft picks may be able to command amounts towards the maximum for the pick position and may be able to negotiate more guaranteed amounts in their contract. However, the scope for negotiation is small because drafted players are unable to refuse the contract and play for a rival franchise. Additionally, the collective bargaining agreement imposes further restrictions on spending on rookie wages in addition to the total franchise salary cap. There is also a minimum spend for the Franchise; at least 90% of the salary cap. The result is a very tight correlation between the draft pick rank order and player wages for the first four years of the draft class.

Players who are unsuccessful in the NFL draft, can still be picked up as undrafted free agents. Such players are paid significantly less than the last picks in the draft. In 2022, 7th round draftees earned a minimum of \$705,000, while the minimum for an undrafted player is \$160,000.

Neither drafted players or undrafted players are guaranteed to play. However, our data suggest drafted players are very likely to play at least one NFL game (>90%). Once a draftee has signed to a Franchise they will join the preseason training camps along with any undrafted players that the Franchise has picked up. During preseason, the 90 man squad (including last season's players) is reduced to the playing roster of 53 players by 30th August (the season begins in September). Players who do not make the roster are either released or kept by the Franchise and assigned to their 'Practice Squad'. Cut players keep their signing bonus (which is guaranteed) but do not get their salary which is usually conditional upon making the roster. Practice squad players will get a fraction of the roster salary. Players can be brought up from the Practice Squad during the season to cover injuries, or promoted from the Practice Squad to the playing roster for the following season.