RESPONDING TO ANTICIPATED BIAS: EVIDENCE FROM FOOTBALL IN ISRAEL

Noam Binstok and Asaf Zussman*
Department of Economics
The Hebrew University of Jerusalem

August 17, 2023

Abstract

How does the anticipation of bias change the behavior of affected agents? This paper provides evidence on this question by studying the decisions of referees and coaches in football (soccer) leagues in Israel over 15 years. We first provide robust evidence that referees exhibit ethnic (Arab/Jewish) in-group bias in penalty card decisions, which affects aggregate game outcomes. We then show that coaches respond to anticipated bias by increasing the number of players matching the ethnicity of the referee in the starting lineups of their teams and that the response is stronger when the anticipated bias is larger.

JEL classification codes: D84, D91, J15, J71, M50, Z22
Keywords: Bias, Discrimination, Ethnicity, Football, Israel.

* E-mails: noam.binstok@mail.huji.ac.il and asaf.zussman@mail.huji.ac.il. We thank Oren Rigbi and participants of presentations at the Israel Economic Association meeting 2023, the Hebrew University of Jerusalem, the Society of Labor Economics meeting 2023, and the University of Haifa for valuable comments.

1. INTRODUCTION

In many social situations, agents may anticipate being adversely affected by bias. How do agents respond to anticipated bias? This paper provides empirical evidence on this question by studying the decisions of Arab and Jewish referees and coaches in Israeli football (soccer) leagues for men. The context offers several advantages for studying bias and the response to anticipated bias. First, the assignment of referees to games is quasi-random. This facilitates the credible identification of causal effects. Second, there is a lot of data. We have information for the universe of games played in these leagues during 15 seasons (from the 2006-2007 season to the 2020-21 season), yielding hundreds of thousands of observations at the player-game level. This allows us to obtain precise estimates. Third, the context facilitates the classification of ethnicity not only at the individual level but also uniquely at the team level. As we explain below, this enables us to study how ethnic bias and the response to anticipated ethnic bias change with the salience of ethnicity at the game level.

Our analysis of referee bias focuses on penalty card decisions and the 11 players who make up the starting lineup of each team participating in a game. According to football rules, a referee may show a player a yellow card to caution him in certain circumstances, e.g., when the player displays unsportsmanlike behavior, argues with the referee, or delays the game. A player warned with a yellow card twice during a game is automatically shown a red card and ejected. In certain circumstances, e.g., when a player commits a severe foul, acts violently against the referee, or uses his hand to stop a goal (when not the goalkeeper), a referee may show the player a red card even without cautioning him first. The analysis focuses on penalty card decisions for two reasons. First, unlike in some other routine decisions, referees enjoy much discretion in making penalty card decisions, which opens the door for bias. Second, referee decisions about penalty cards may significantly affect game outcomes.

Analysis of the data at the player-game level yields robust evidence of ethnic in-group bias – the tendency to favor members of one's ethnic group – in referee penalty card decisions. Estimated bias is not only statistically significant but also economically

meaningful. Our preferred specification shows that the likelihood that a player in the starting lineup of his team will be shown a yellow card is 1.0 percentage points – or 5.7 percent of the mean – lower when the player is from the same ethnic group as the referee. The likelihood that a player in the starting lineup of his team will be shown a red card is 0.3 percentage points – or 21.0 percent of the mean – lower when the player is from the same ethnic group as the referee.

A well-established finding in the literature on social identification in general and in-group bias in specific is that the strength of these phenomena tends to increase when group membership is salient. The context we study offers a unique opportunity to study the effect of the salience of group membership. Because of the extreme level of residential ethnic segregation in Israel and the fact that teams represent specific localities (or neighborhoods), teams are essentially identified as either Arab or Jewish (e.g., Maccabi Akhi Nazareth represents the Arab city of Nazareth while Maccabi Petach-Tikva represents the Jewish city of Petach-Tikva). This is so even though the starting lineups of Arab teams often include Jewish players and vice versa (we provide evidence of these patterns below). The context thus naturally generates situations where the salience of ethnicity is particularly high: cross-ethnicity games, i.e., games in which one team is Arab, and the other is Jewish. In the data we analyze, cross-ethnicity games constitute almost half of the total.

Our analysis at the player-game level shows that ethnic in-group bias in referee penalty card decisions is much stronger in cross-ethnicity games than in same-ethnicity games (in which both teams are Arab or Jewish). In cross-ethnicity games, the likelihood that a player will be shown a yellow card is 1.7 percentage points – or 9.6 percent of the mean – lower when the player is from the same ethnic group as the referee. The likelihood that a player will be shown a red card is 0.7 percentage points – or an astonishing 43.2 percent of the mean – lower when the player is from the same ethnic group as the referee. Both results are highly significant. The corresponding effects for same-ethnicity games are an order of magnitude smaller and statistically insignificant.

Aggregating the data to the team-game level, we find that referee ethnic in-group bias has a meaningful effect on the number of yellow cards, red cards, goals scored, and points awarded to the team. This analysis takes into account the fact that the ethnic composition

of the starting lineup of the team affects referee decisions in two ways: for instance, an increase in the number of starters who match the ethnicity of the referee not only reduces the number of yellow cards shown by the referee to the team's players (direct effect) but also increases the number of yellow cards shown by the referee to its opponent's players (indirect effect) – the total effect factors in these two effects. Relative to the mean of the dependent variable, the estimated total effect of referee ethnic in-group bias amounts to 7.1 percent for yellow cards, 5.5 percent for red cards, 1.3 percent for goals, and 1.8 percent for points. All these effects are highly statistically significant. Comparing the results between cross-ethnicity games and same-ethnicity games, we find that except for yellow cards, the effect is stronger and more statistically significant in cross-ethnicity games.

The results presented above suggest that coaches might want to take into account the ethnic identity of the referee in each game. In particular, bias incentivizes coaches to select more Arab players from their teams' rosters to be members of the starting lineups when the referee is Arab and to do the opposite when the referee is Jewish. This adjustment is possible because lineup decisions can be made even on the game day while the referee's identity is known well in advance.

Analyses of starting lineup decisions at both the team-game level and the player-game level show that coaches indeed respond to the ethnic identity of the assigned referee in the hypothesized direction. To illustrate, the likelihood that a coach will select a player from his team's roster to be a member of the starting lineup increases by 0.95 percentage points (1.5 percent of the mean) when the player's ethnicity matches that of the referee.

Each of the 11 players in a football team's starting lineup is assigned to a particular position. One goalkeeper and ten outfield players fill various defensive, midfield, and attacking positions. These positions describe the player's primary role and area of operation on the field. The fact that players are not perfectly substitutable with others implies that adjusting the ethnic composition of the starting lineup in response to referee ethnicity is costly. The adjustment is arguably especially challenging when the number of Arab players in the team's roster is very low and when the number of Arab players in the team's roster is very high (or, equivalently, the number of Jewish players is very low). Our analysis provides evidence consistent with this hypothesis. We find that the strength of lineup

adjustment depends on the number of Arab players on the team roster: the effect is particularly large and statistically significant when this number is in an intermediate range.

Another likely key consideration in adjusting the starting lineup is player quality. Intuitively, it seems reasonable to assume that coaches will be less likely to involve high-quality players in the adjustment. To study this issue, we create a proxy for quality: the mean number of goals scored by the player in games he played as a starter. Consistent with the hypothesis, our analysis finds that lineup adjustment is performed using relatively low-quality players (non-scorers) while high-quality players (scorers) are excluded.

An important question is whether coaches adjust the ethnic composition of the starting lineup merely based on a simple rule of thumb – increase the number of starters who match the ethnicity of the referee – or instead use a more sophisticated thought process based on the anticipated extent of referee bias. We provide two pieces of evidence in support of the latter alternative.

First, we show that the lineup adjustment in response to the ethnic identity of the assigned referee is much more substantial and statistically significant in cross-ethnicity games than in same-ethnicity games. This is consistent with our previous finding that in-group bias in referee decisions is stronger in cross-ethnicity games.

Second, we use a two-step procedure to show that coaches respond to referee-specific bias levels. In the first step, we measure the extent of ethnic in-group bias at the referee level. In the second step, we estimate the effect of the measured bias of the assigned referee on starting lineup decisions. The analysis yields strong evidence that coaches respond to the extent of anticipated referee-specific ethnic in-group bias: the effect on starting lineups increases monotonically as we progressively focus on the more biased referees from each ethnic group. The estimated effect on starting lineups for the 20 percent most biased referees from each ethnic group is about twice the size of that estimated based solely on referee ethnicity (i.e., using the rule-of-thumb).

This paper is related to two strands of literature. The first is the literature on in-group bias, which is part of a vast body of research on discrimination.[1] This literature has studied in-group bias based on various identity markers, such as race, ethnicity, and gender, in decisions in different market and non-market contexts, such as hiring, lending, policing, judicial sentencing, grading, and more.[2] Within this literature, a few papers study racial and ethnic in-group bias in sports refereeing. The two most prominent studies along these lines are Price and Wolfers (2010) and Parsons et al. (2011), who provide evidence of racial and ethnic in-group bias in, respectively, professional basketball (NBA) and baseball (MLB) in the United States.[3] We contribute to the literature on in-group bias (and more generally to the literature on discrimination) by providing stark evidence that the salience of group membership has a crucial effect on the strength of bias.

It has long been recognized theoretically that anticipation of bias changes agents' behavior. See, for example, Welch (1967), Arrow (1973), Lundberg and Startz (1983), Coate and Loury (1993), and Knowles, Persico, and Todd (2001). However, empirical evidence on such behavior, especially from outside the lab, is relatively rare.[4] Thus, probably the most significant contribution of our paper is to provide well-identified real-world evidence on how the anticipation of bias changes the behavior of affected agents.

The rest of the paper is organized as follows. Section 2 provides information on the context and the data. Section 3 provides evidence on referee in-group bias and the effect of

[1] Bertrand and Duflo (2017), Neumark (2018), and Lang and Kahn-Lang Spitzer (2020) provide recent reviews of the literature on discrimination.
[2] Shayo (2020) provides a recent review of the literature on in-group bias.
[3] Price and Wolfers (2010) use data from the 1991-92 to the 2003-04 NBA seasons. Using data for later seasons, Pope, Price, and Wolfers (2018) and Mocan and Osborne-Christenson (2022) report finding no overall racial in-group bias in refereeing in the NBA.
[4] Parsons et al. (2011) show that pitchers throw pitches that allow umpires less subjective judgment (e.g., fastballs over home plate) when anticipating umpire bias. Zussman (2013), who studies the Israeli online market for used cars, shows that Arab sellers, subject to customer discrimination in this market, manipulate their ethnic identity by leaving the name field in their advertisements blank. Examining the performance of cashiers in a French grocery store chain, Glover, Pallais, and Pariente (2017) find that when minority cashiers, but not majority cashiers, are scheduled to work with ethnically biased managers (as determined by an implicit association test), they are absent more often, spend less time at work, scan items more slowly, and take more time between customers. Kudashvili and Lergetporer (2022) provide evidence from a lab experiment on the strategic response to anticipated discrimination.

the salience of ethnicity on it. Section 4 shows how coaches adjust their teams' starting lineups in anticipation of referee bias. Section 5 concludes.

2. CONTEXT AND DATA

2.1. Football leagues and referees

The Israel Football Association runs club football for men in Israel. The system consists of five interconnected sets of leagues, organized in a hierarchical structure: based on their performance, teams can be promoted from a lower group of leagues to a higher one or relegated from a higher group of leagues to a lower one. The top leagues are the Premier League (the highest league in the system) and the National League, which comprise7 14 and 16 clubs, respectively. These are professional leagues that operate at the national level. The A and B leagues are semi-professional and operate at a regional level. The A-League comprises 32 member clubs and is divided into two regional leagues. The B League comprises 64 member clubs and is divided into four regional leagues. The C League, which is for amateurs, consists of around 110 member clubs (the number of clubs changes over time) and is divided into eight regional leagues. Thus, the system consists of 16 leagues and around 240 teams.[5]

To become a football referee in Israel, candidates must pass a five-day refereeing course run by the Referee Union, an organization related to the Israel Football Association. The course covers the rules of the game and their enforcement. At the end of the course, candidates must pass a test on the study materials and a fitness test. Candidates who pass these tests are mentored for a short period by tenured colleagues and then receive a refereeing certificate and can start working. New referees first work in leagues for children and adolescents and then in amateur leagues for men. There are typically three referees in the games we study: a main referee and two assistant referees (linesmen). The main referee is solely responsible for all major decisions, including those concerning penalty cards.

[5] Until the 2008-09 season, the system consisted of six sets of leagues. The additional league was positioned between the National League and the A League. Our analysis takes this change into account.

2.2. Data

The data used in this paper were scraped from the website of the Israel Football Association. They cover the universe of games played in all leagues between the 2006-07 season – the first for which data are available on the website – and the 2020-21 season. The website reports the summary of each game in a structured ("box score") format. This includes the following information for each match: league, season, round[6], week of play[7], date, start time, final score, score at halftime, location of the stadium, name of the home team, name of the away team, list of goals scored (for each: full name of the scorer and minute), players in the starting lineup of each team (for each player: full name and number, minute substituted out, type of penalty card/s shown and minute/s), players who were substituted in by each team (for each player: full name and number, minute substituted in, type of penalty card/s shown and minute/s)[8], other players in the team's roster who did not participate in the game (for each player: full name and number)[9], full name of the coach of each team, and full names of the members of the referee crew (main referee and assistant referees).

The ethnicities of players and referees play a critical role in the analysis. To classify ethnicity, we rely on a name-based procedure similar to that used by Shayo and Zussman (2011). The procedure builds on the fact that Arabs and Jews have very different naming conventions and utilizes a dataset derived from the Israeli Population Registry, which provides, separately for each gender, the probabilities that specific first and last names belong to an Arab citizen of Israel. We first classify first and last names as Arab (Jewish) if they are at least twice as prevalent among Arabs (Jews) than among Jews (Arabs). We then classify the full name as Arab (Jewish) if both the first name and the last name are classified as Arab (Jewish), if the first name is classified as Arab (Jewish) and the last name is not classified as either Arab or Jewish, or if the last name is classified as Arab (Jewish)

[6] In all leagues, each team plays against all other teams in the league it belongs to twice, once at home and once away from home, in two rounds of matches. In addition, some teams may play in a playoff round.
[7] The number of weeks of play in a season depends on the number of teams in the league and on the playoff structure.
[8] Teams are allowed up to five substitutions in a game. Before the 2019-20 season, the two professional leagues allowed only three substitutions. A player who was substituted out cannot be substituted back in.
[9] Team rosters include up to 20 players.

and the first name is not classified as either Arab or Jewish. Using this procedure, we classify the ethnicity of 96 percent of the players and 100 percent of the referees.[10] Our analysis excludes observations for which we could not classify the ethnicity of either the player or the referee.

2.3. Referee assignment

The ability to credibly identify bias depends on the assumption that the assignment of referees to games is as good as random. Referees in all leagues are assigned to games by the Assignment Committee of the Referee Union. There is no publicly available official information on the exact procedures used by this committee. However, informal conversations with referees as well as publicly available information (e.g., from interviews with committee members) suggest that the committee factors in several considerations in assigning referees to games, including referee seniority, availability, and geographical constraints; the ethnic makeup of the teams participating in a game is not taken into account.

Figure 1 provides evidence confirming that referee assignment is effectively random. Each point in the figure represents data for a specific league and season. The figure plots the actual vs. the predicted share of player-game observations in which a player on the rosters of the teams taking part in the game and the referee officiating the game are Arab. The predicted share is calculated from the share of players (in teams' rosters in all games) who are Arab and the share of games officiated by Arab referees by assuming random matching. If referee assignment is effectively random, we would expect the actual share to be close to the predicted share. The pattern observed in Figure 1 is consistent with a quasi-random assignment: with a few exceptions, the points lie very close to the upward-sloping 45-degree line.[11]

[10] In addition, we classify the ethnicity of 89 percent of the coaches. This relatively low classification success rate is due to the fact that, in many cases, the "box score" does not report coaches' names.

[11] Figure A1 in the Appendix shows the results of a similar graphical test of quasi-random assignment of referees to games. It focuses on ethnic identity at the team rather than the player level. This approach builds on the fact, noted in the introduction, that in the Israeli context, it is possible to classify team ethnicity. The evidence further supports the claim that referee assignment is effectively random.

[Figure 1]

2.4. Summary statistics

Table 1 provides summary statistics at the player-game level, first for all players and then by player ethnicity. The analysis is restricted to players who were identified as either Arab or Jewish and included in the starting lineups of their teams. The analysis is further limited to games with at least one Arab player and at least one Jewish player in total in the starting lineups (we can estimate ethnic in-group bias only in such games). On average, players receive 0.173 yellow cards and 0.015 red cards and score 0.127 goals per game. The likelihood that a player will face an Arab (main) referee is 13.5 percent; the corresponding figure for the assistant referees is around 20.5 percent. The shares of observations associated with Arab teams, Arab coaches, and cross-ethnicity games are 42.2, 32.2, and 43.8 percent, respectively.

[Table 1]

Arab players account for 40.4 percent of the observations. Table 1 indicates that relative to Jewish players, Arab players receive more yellow and red cards and score slightly fewer goals per game. Arab players are more likely than Jewish players to face Arab referees. This is explained by the fact that Arab and Jews tend to live in different parts of the country – in particular, the highest concentration of Arabs is in northern Israel – and that most leagues are organized on a geographical basis. The table also shows that Arabs are much more likely than Jews to play for Arab teams and to be coached by Arab coaches. Finally, Arab players are more likely than Jewish ones to play in cross-ethnicity games.

3. ETHNIC IN-GROUP BIAS IN REFEREE DECISIONS

3.1. Evidence of bias at the player-game level

Figure 2 uses raw player-game level data to illustrate ethnic in-group bias in referee penalty card decisions. Panel A focuses on yellow cards. It shows the mean number of yellow cards per game for each referee and player ethnicity combination. The two left columns focus on decisions made by Arab referees. Arab referees show, on average, 0.174

yellow cards to Arab players and 0.166 yellow cards to Jewish players, a difference of 0.008 cards per game. The right two columns in panel A focus on decisions made by Jewish referees. These referees show, on average, 0.184 yellow cards to Arab players and 0.166 yellow cards to Jewish players, a difference of 0.018 cards per game. The difference between these differences, equal to -0.010, is a raw measure of ethnic in-group bias in referee yellow card decisions: it shows by how much referees are less likely to show yellow cards to players from their ethnic group relative to players from the other ethnic group. A similar exercise performed for red cards (Panel B) shows that the raw measure of ethnic in-group bias is, in this case, -0.002.

It is important to note that, without an objective measure of "correct" decisions about yellow and red cards, it is impossible to determine which type of referee is biased. It is possible that only Arab referees are biased, only Jewish referees are biased, or both Arab and Jewish referees are biased. It is also impossible to determine whether bias takes the form of endophilia (favoring members of one's group) or exophobia (discriminating against members of the other group).[12] Thus there are many possible interpretations of the patterns observed in Figure 2. However, some interpretations – e.g., that both Arab referees and Jewish referees favor members of their ethnic group – intuitively seem to be more likely than others. [13]

We next analyze ethnic in-group bias in referee penalty card decisions econometrically. We start with a standard baseline specification (see Bar and Zussman, 2020):

$$PenaltyCard_{ijklsr} = \alpha_0 + \alpha_1 ArabPlayer_i + \alpha_2 ArabReferee_r + \alpha_3 ArabPlayer_i * ArabReferee_r + \varepsilon_{ijklsr} \quad (1)$$

[12] See Feld, Salamanca, and Hamermesh (2016) on this distinction.

[13] A possible concern about our analysis, common to the literature studying in-group bias in natural settings, is that the empirical patterns observed in Figure 2 may reflect players' endogenous reactions to the assigned referee's ethnicity rather than referee bias. Specifically, the argument is that players may become more aggressive and receive more yellow and red cards when the assigned referee is from the out-group. We do not have data that would allow us to rule out this alternative interpretation. However, it is likely invalid. Awareness of possible referee bias should make players behave *less* aggressively when the assigned referee is from the out-group. Such behavior, consistent with the behavior of coaches as we document below, would imply that our estimate of referee bias is a *lower bound*.

where $PenaltyCard_{ijklsr}$ is an indicator variable that takes the value of one when player *i*, playing for team *j*, against team *k*, in a game held in location *l* (home or away for team *j*), in season *s*, and refereed by referee *r*, was shown a specific type of penalty card (either yellow or red) during the game; $ArabPlayer_i$ is an indicator variable taking the value of one for Arab players; $ArabReferee_r$ is an indicator variable taking the value of one for Arab referees; and ε_{ijklsr} is an error term clustered within the combination of referee and season.

Equation (1) allows for possible differences between Arab and Jewish players and between Arab and Jewish referees that should not be interpreted as evidence of bias. First, Arab and Jewish players may have different characteristics, such as field positions, making them more or less likely to be shown penalty cards. The coefficient α_1 captures this. Second, Arab and Jewish referees may have different characteristics, e.g., refereeing style, that make them more or less likely to show penalty cards. The coefficient α_2 captures this. Interest focuses on the coefficient α_3. If $\alpha_3 < 0$ referees exhibit ethnic in-group bias, i.e., they display relative favoritism toward players from their ethnic group.

Estimating equation (1) yields evidence of in-group bias in referee decisions about yellow and red cards (Table 2). The likelihood that a player will be shown a yellow card is 1.02 percentage points – or 5.9 percent of the mean of the dependent variable – lower when the player is from the same ethnic group as the referee (column 1). The likelihood that a player will be shown a red card is 0.21 percentage points – or 14.3 percent of the mean – lower when the player is from the same ethnic group as the referee (column 5).

[Table 2]

We next gradually augment equation (1) with a set of controls. The full specification takes the following form:

$$\begin{aligned} PenaltyCard_{ijklsr} \\ &= \alpha_0 + \alpha_3 ArabPlayer_i * ArabReferee_r + \beta_{ils} + \theta_{rs} + \gamma_{jls} \\ &+ \delta_{kls} + \varepsilon_{ijklsr} \end{aligned} \quad (2)$$

where β_{ils} is a fixed effect for the combination of player, game location (home or away), and season; θ_{rs} is a fixed effect for the combination of referee and season; γ_{jls} is a fixed effect for the combination of player i's team (team j), game location, and season; δ_{kls} is a fixed effect for the combination of the other team, game location, and season; and the other variables are defined as above. The β_{ils} fixed effect captures the player's propensity to be shown penalty cards (which, in turn, may reflect player characteristics such as field position), allowing this propensity to vary by game location and season. The θ_{rs} fixed effect captures the tendency of the referee to show penalty cards, enabling this tendency to vary by season. The γ_{jls} and δ_{kls} fixed effects capture possible influences on player i's likelihood of being shown a penalty card emerging from his team and from the other team, where these influences are allowed to vary by game location and season.[14]

Estimating equation (2) continues to yield evidence of in-group bias in referee penalty card decisions. Note that while adding the control variables leads to a dramatic increase in the explanatory power of the regression, it has relatively little effect on the size of the main coefficient of interest (this is especially true for yellow cards). This provides additional support for the claim that the assignment of referees to games is quasi-random. Under the full specification of equation (2), the likelihood that a player will be shown a yellow card is 1.00 percentage points – or 5.7 percent of the mean – lower when the player is from the same ethnic group as the referee (column 4). The likelihood that a player will be shown a red card is 0.31 percentage points – or 21.1 percent of the mean – lower when the player is from the same ethnic group as the referee (column 8).[15]

[14] Equation (2) excludes the variables $ArabPlayer_i$ and $ArabReferee_r$ because it includes player and referee fixed effects.

[15] Tables A1-A3 in the Appendix use equation (2) to explore heterogeneity in referee bias by, respectively, type of league, stage in the season, and type of referee crew. Table A1 shows that referee bias is much stronger in the amateur leagues (C leagues) than in the professional and semi-professional leagues (the premier, national, A, and B leagues). This difference could be due to various reasons. Among them is that professional and semi-professional league referees are more experienced and subjected to more oversight. Table A2 shows that referee bias is much stronger in the last quarter of the season than in the first three quarters. This difference may reflect the influence of the higher stakes in the latter part of the season. Table A3 addresses the following question: Does the extent of in-group bias exhibited by the main referee depend on the ethnicities of the two assistant referees? A natural hypothesis is that it will be easier for the main referee to exhibit bias if the other two referees share his ethnicity than if at least one is from the other ethnic group. To test this hypothesis, we compare the extent of bias between games with "homogeneous" referee crews and games with "nonhomogeneous" crews. The evidence is mixed: homogeneous crews exhibit much stronger ethnic in-group bias than nonhomogeneous crews in the case of yellow cards, but this is not true in the case

3.2. The effect of the salience of ethnicity on bias

A large body of research, primarily based on lab experiments, shows that in-group bias tends to increase when group membership is salient. We leverage a unique feature of the context to study whether this effect is present in the football games studied here. Israel is characterized by an extreme level of residential ethnic segregation: the population of the vast majority of localities in the country is essentially either all Jewish or all Arab; the eight localities that are officially considered "integrated" – according to the definition used by the Israeli Central Bureau of Statistics, a locality is considered integrated if the share of Arabs in the total population is between 2 and 50 percent – are ethnically segregated by neighborhoods.

Since football teams represent specific localities (or neighborhoods), they are effectively identified with a particular ethnicity. This is so even though many Arab teams have Jewish players and vice versa. Figure 3 illustrates this point by showing histograms of the number of Jewish starters in Arab teams (Panel A) and the number of Arab starters in Jewish teams (Panel B). Arab teams have at least one Jewish starter in about half of their games, and Jewish teams have at least one Arab starter in about a third of their games.

[Figure 3]

The situation described above thus implies that the games analyzed in this paper can be classified as either cross-ethnicity – when one team is Arab, and the other is Jewish – or same-ethnicity – when both teams are Arab, or both teams are Jewish. In comparing outcomes across the two types of games, we assume that the salience of ethnicity is higher in cross-ethnicity games. This assumption is supported by ample anecdotal evidence. An extreme illustration comes from Beitar Jerusalem Football Club, one of the most popular teams in Israel, whose fan base is notorious for its anti-Arab and anti-Muslim attitudes. Beitar is the only club in the Israeli Premier League to have never signed an Arab player.[16]

of red cards. For a similar analysis in the context of refereeing in international cricket matches, see Fernando and George (forthcoming).

[16] Midway through the 2012–13 season, Beitar's owner brought two Muslim players from Chechnya to play for the team, which led fans to conduct a violent campaign to reverse the decision. These events are portrayed in the Emmy-winning documentary *Forever Pure*.

Games between Beitar and Arab teams often result in violent confrontations between fans of the opposing sides.

To examine whether the salience of ethnicity affects the extent of referee ethnic in-group bias, we estimate equation (2) separately for cross-ethnicity and same-ethnicity games. Results are reported in Table 3. The ethnic in-group bias in referee decisions documented in Table 2 for all games is mainly driven by cross-ethnicity games. In such games, the likelihood that a player will be shown a yellow card is 1.66 percentage points – or 9.6 percent of the mean – lower when the player is from the same ethnic group as the referee (column 1). The corresponding estimate for same-ethnicity games is 0.17 percentage points and statistically insignificant. Shifting our focus to red cards, we find that the likelihood that a player will be shown such a card is 0.67 percentage points – or an astounding 43.2 percent of the mean – lower when the player is from the same ethnic group as the referee (column 3). The corresponding estimate for same-ethnicity games is 0.03 percentage points and statistically insignificant.

[Table 3]

3.3. Evidence of bias at the team-game level

The analysis above provides evidence at the player-game level of referee ethnic in-group bias in penalty card decisions. A natural question follows: how does referee ethnic bias – in penalty card decisions and potentially in other decisions[17] – affect outcomes at the team-game level?

When studying this question, it is essential to realize that the analysis at the player-game level may miss an important aspect of bias. Consider a game between team *j* and team *k*. All else being equal, an increase in the number of team *j* players who match the ethnicity of the referee will not only lead a biased referee to show *fewer* penalty cards to team *j* players but also to show *more* penalty cards to team *k* players. Price and Wolfers (2010) call the former effect the *direct effect* and the latter effect the *indirect effect*. We are

[17] Below we discuss referee decisions concerning 11m penalty kicks. Another type of referee decision that bias may influence is stoppage time. Garicano, Palacios-Huerta, and Prendergast (2005) report evidence of referee bias favoring home teams in stoppage times decisions. We do not have data on stoppage time.

interested in estimating the *total effect*, which is the difference between the two effects, i.e., team *j* benefits from having fewer cards shown to its players and more cards shown to k's players. This rationale naturally carries over to other outcomes of interest.

Figure 4 illustrates ethnic in-group bias in referee yellow card and red card decisions at the game level. The figure captures the concept of the *total effect* by focusing on *differences* in the numbers of yellow cards, red cards, and Arab starters between the two teams taking part in the game. Panel A shows that in matches with Arab referees, the mean difference in the number of yellow cards between the home and away teams is positively associated with the difference in the number of Arab starters between the home and away teams. Panel B shows a similar association for games with Jewish referees. However, the association is much stronger for matches with Jewish referees than for games with Arab referees. Similar patterns are observed for red cards (Panels C and D). These patterns suggest the existence of in-group bias in yellow and red card decisions.

[Figure 4]

Following Price and Wolfers (2010), we conduct the econometric analysis of the effect of referee bias on game-level outcomes using a dataset with two observations per game. To illustrate, for a match between team *j* and team *k*, in one observation team *j* will be the main team and team *k* its opponent, and in the other observation team *k* will be the main team and team *j* its opponent. This allows us to estimate the direct, indirect, and total effects of bias on team outcomes. We estimate the following equation:

$$\begin{aligned} GameOutcome_{tolsr} = \alpha_0 + \alpha_1 ArabStarters_{tls} + \alpha_2 ArabStarters_{tls} * \\ ArabReferee_r + \alpha_3 ArabStarters_{ols} + \alpha_4 ArabStarters_{ols} * \\ ArabReferee_r + \theta_{rs} + \gamma_{tls} + \delta_{ols} + \varepsilon_{tolsr} \end{aligned} \quad (3)$$

where $GameOutcome_{tolsr}$ is the outcome of a game held between team *t* and its opponent *o*, in location *l*, in season *s*, and refereed by referee *r*; $ArabStarters_{tls}$ and $ArabStarters_{ols}$ are, respectively, the number of Arab players in the starting lineups of team *t* and *o* in that game; $ArabReferee_r$ is an indicator variable taking the value of one for Arab referees; θ_{rs} is a fixed effect for the combination of referee and season; γ_{tls} is a fixed effect for the combination of team, game location (home or away), and season for

team t; δ_{ols} is a fixed effect for the combination of the team, game location, and season for team o; and ε_{tolsr} is an error term clustered within the combination of referee and season. In this equation, the coefficient α_2 captures the direct effect, the coefficient α_4 captures the indirect effect, and the difference between the two coefficients, $\alpha_2 - \alpha_4$, captures the total effect. We study the effect of bias on four outcome variables: yellow cards, red cards, goals scored, and points awarded (zero points are awarded for a loss, one for a draw, and three for a victory). Results are reported in Table 4.

[Table 4]

Focusing first on yellow and red cards, we find that the direct effects estimated using the team-game level data (columns 1 and 2, first row) are somewhat larger in magnitude than the (direct) effects estimated using the individual level data (Table 2). The indirect effects (second row) are similar in absolute size to the direct effects. The total effects (middle panel) are not only statistically significant but also economically significant. For yellow cards, the total effect implies that *one* additional player who matches the ethnicity of the referee in the starting lineup of the team tilts the balance of yellow cards by 0.028, equivalent to 1.3 percent of the mean of the dependent variable (which, it is important to stress, is measured at the team level, i.e., for 11 starters). Similarly, for red cards, the total effect implies that *one* additional player who matches the referee's ethnicity in the team's starting lineup tilts the balance of cards by 0.010, equivalent to 5.5 percent of the mean of the dependent variable.

Given the strength of ethnic in-group bias in referee decisions we document for yellow and red cards, it is not surprising that we also estimate significant total effects for other, arguably more important, game outcomes (columns 3 and 4). We find that *one* additional player who matches the referee's ethnicity in the team's starting lineup tilts the balance of goals and points by 1.3 and 1.8 percent of the mean of the dependent variable.[18]

[18] Decisions concerning 11m penalty kicks are highly consequential for game outcomes and involve much discretion from the referee. Such kicks are awarded when an offense punishable by a direct free kick is committed by a player in his penalty area (i.e., in the vicinity of his own goal). A player performs the kick from the 11m mark, while the goal is defended only by the opposing team's goalkeeper. We do not have data on referee decisions on 11m penalty kicks, but we do have data on goals scored in 11m penalty kicks (penalty kicks are usually successful; for example, using data from the elite French and Italian leagues, Chiappori,

Table 5 reports results from estimating equation (2) separately for cross-ethnicity and same-ethnicity games. For cross-ethnicity games, we find that the total effects of bias have the expected signs and are large and statistically significant. For same-ethnicity games, we find such a pattern only for yellow cards; the estimated effects for the three other outcome variables are relatively small and statistically insignificant.

[Table 5]

4. STARTING LINEUP DECISIONS

4.1. Response to referee ethnicity

The results presented in the previous section suggest that coaches might want to take into account the ethnic identity of the referee in each game. In particular, referee ethnic in-group bias incentivizes coaches to include more Arab players in their teams' starting lineups when the referee is Arab and to do the opposite when the referee is Jewish. This section studies this hypothesized starting lineup adjustment.

We begin the investigation by analyzing data at the team-game level. The analysis is restricted to teams that have on their rosters at least one Arab player and at least one Jewish player since only in such cases can the adjustment of the ethnic composition of the starting lineup take place. Figure 5 uses raw data to present evidence on starting lineup adjustment. For each number of Arab players in the team roster and each referee ethnicity, it plots the mean (and the associated 95% confidence interval) of the ratio between the number of Arab starters and the number of Arab players on the roster.

[Figure 5]

The figure presents an interesting pattern. When the number of Arab players on the roster is low, the ratio is slightly higher when the referee is Arab rather than Jewish. When the number of Arab players on the roster is high – or, equivalently, when the number of

Levitt, and Groseclose (2002) report a success rate of 75 percent). We thus estimate equation (3) for penalty kicks scored. Results are reported in Table A4 in the Appendix. The estimated total effect implies that *one* additional player who matches the referee's ethnicity in the team's starting lineup tilts the balance of goals scored in a 11m penalty kick by 4.2 percent of the mean of the dependent variable. This effect is more than three times larger than the one estimated in column 3 of Table 4 for the total number of goals.

Jewish players is low – the ratio does not seem to be associated with the ethnicity of the assigned referee. In contrast, when the number of Arab players on the roster is in an intermediate range, the ethnic identity of the referee matters a lot: the share of Arab players on the roster of the team chosen to be members of the starting lineup is significantly higher when the referee is Arab rather than Jewish. The observed pattern suggests that adjusting the ethnic composition of the starting lineup in response to referee ethnicity is less costly when the roster includes sufficiently large numbers of Arab and Jewish players.

We study the question of starting lineup adjustment at the team-game level econometrically by estimating the following equation:

$$ArabRatio_{jklsr} = \alpha_0 + \alpha_1 ArabReferee_r + \gamma_{jls} + \delta_{kls} + \varepsilon_{jklsr} \qquad (4)$$

where $ArabRatio_{jklsr}$ is the ratio between the number of Arab starters and the number of Arab players on the roster of team *j*, playing against team *k*, in a game held in location *l* (home or away for team *j*), in season *s*, and refereed by referee *r*; all the other variables are defined as above. Results are presented in Table 6. We are interested in the coefficient α_1, which measures the effect of referee ethnic identity on the ethnic composition of the starting lineup.

[Table 6]

The overall effect of referee ethnicity on the ratio between the number of starters and the number of Arab players on the roster is estimated at 1.1 percentage points, or about 1.8 percent of the mean of the dependent variable, and is statistically significant (column 1). When we break down the analysis by the number of Arab players on the roster (columns 2-4), the results echo those shown in Figure 5: the effect is particularly large and statistically significant when the number of Arab players is in an intermediate range: the effect is estimated at 1.5 percentage points or 2.4 percent of the mean of the dependent variable.

We next study the effect of referee ethnicity on starting lineup decisions using data at the player-game level, estimating the following equation:[19]

$$Starter_{ijklsr} = \alpha_0 + \alpha_1 ArabPlayer_i * ArabReferee_r + \theta_{rs} + \beta_{ils} + \gamma_{jls} + \delta_{kls} + \varepsilon_{ijklsr} \quad (5)$$

where $Starter_{ijklsr}$ is an indicator variable that takes the value of one when player *i*, belonging to the roster of team *j*, playing against team *k*, in a game held in location *l* (home or away for team *j*), in season *s*, and refereed by referee *r*, was in the starting lineup of his team; all the other variables are defined as above.

The results presented in column 1 of Table 7 indicate again that coaches adjust the ethnic composition of their teams' starting lineups in response to the ethnic identity of the assigned referee. The likelihood that a coach will include a player of a given ethnicity (instead of a player from the other ethnic group) in the starting lineup of his team increases by 0.95 percentage points, or 1.5 percent of the mean of the dependent variable, when the player is from the same ethnic group as the referee. When we break down the analysis by the number of Arab players on the roster, the results again suggest that the effect is particularly large and statistically significant in the intermediate range. In this range, the estimated effect is 1.7 percentage points or 2.6 percent of the mean of the dependent variable.

[Table 7]

Another important consideration in the adjustment of the starting lineup should be player quality. Coaches will likely tend to avoid involving high-quality players in the adjustment. To study this issue, we created two proxies for player quality. Both are constructed only for players who started in at least 20 games in the period analyzed. The first proxy is based on the mean number of goals scored by the player in games in which he started. We define players in the top quintile of the distribution as "scorers." The second proxy is based on an adjusted measure of the mean number of goals scored. Specifically, we first obtain residuals from a player-game level regression of goals scored on fixed

[19] Appendix Table A5 provides summary statistics at the player-game level for all players, including non-starters.

effects for the combination of referee and season and fixed effects for the combination of team, game location, and season for the two teams participating in the game. We then define players in the top quintile of the distribution of the mean residual as "scorers." We estimate equation (5) for all players in this restricted sample and separately for "non-scorers" and for "scorers." Results are presented in Table 8.

[Table 8]

The estimated effect of referee ethnicity on starting lineup decisions (column 1) is almost identical to that obtained previously (column 1 of Table 7), despite the non-trivial decline in the number of observations available for analysis. More importantly, regardless of which proxy of quality we use, we find that lineup adjustment is carried out using lower-quality players (non-scorers). In contrast, high-quality players (scorers) are excluded from it.

4.2. Response to the strength of anticipated bias

An important question is whether coaches adjust the ethnic composition of the starting lineup merely based on a simple rule of thumb – increase the number of starters who match the ethnicity of the referee – or instead use a more sophisticated thought process that is based on the anticipated strength of referee ethnic in-group bias. We provide two pieces of evidence in support of the latter alternative.

First, since the analysis in Section 2 has shown that in-group bias in referee decisions is stronger in cross-ethnicity games, we expect that the adjustment of the starting lineup will also be stronger in cross-ethnicity games. To study whether this is the case, we rerun the analysis of Table 7 separately for cross-ethnicity and same-ethnicity games. The results, presented in Table 9, are consistent with the hypothesis. When studying all cross-ethnicity and same-ethnicity games, the effect of referee ethnicity on player inclusion in the starting lineup is almost 70 percent stronger in the former type of games than in the latter (columns 1 and 2). The difference in the strength of the effect is starker for games in the intermediate range in terms of the number of Arab players on the roster (columns 5 and 6). In cross-ethnicity games, the likelihood that a coach will include a player of a given ethnicity in the starting lineup of his team increases by 2.8 percentage points (or 4.3 percent of the mean

of the dependent variable) when the player is from the same ethnic group as the referee; the corresponding figure for same-ethnicity games is 1.4 percentage points and statistically insignificant.

[Table 9]

Second, we use a two-step procedure to show that coaches respond to anticipated referee-specific levels of ethnic in-group bias. In the first step, we use the individual level data to measure ethnic bias in refereeing by regressing, for each referee separately, an indicator taking the value of one when the referee shows a player a yellow card on an indicator for Jewish players. These regressions are restricted to players identified as Arab or Jewish and included in their teams' starting lineups and to games where there is at least one Arab player and at least one Jewish player in total in the starting lineups (integrated games). The coefficient on the explanatory variable is a rough measure of (relative) ethnic bias: high values imply that the referee is biased in favor of Arab players, and low values indicate that the referee is biased in favor of Jewish players. To improve precision in the second step of the procedure, we focus on bias estimates obtained for referees who refereed at least 50 integrated games. We concentrate on yellow cards rather than on red cards because yellow cards are shown an order of magnitude more frequently than red cards, which also helps to improve precision.[20]

Figure 6 shows kernel density plots of standardized (z-score) versions of the estimated bias coefficients obtained for 22 Arab and 162 Jewish referees who refereed at least 50 integrated games. As expected, Arab referees exhibit, on average more bias in favor of Arab players than Jewish referees do: the difference in means is about 0.3 of a standard deviation.

[Figure 6]

In the second step of the procedure, we use individual-level data to estimate the effect of measured referee bias on starting lineup decisions. As a preliminary step, we estimate equation (5) for the restricted sample and compare the results (column 1 of Table 10) to

[20] The figures reported in Table 1 indicate that, on average, a player in the starting lineup can expect to be shown a yellow card once every six games and a red card only once every 68 games.

those obtained previously (column 1 of Table 7). Reassuringly, despite a decline of almost 30 percent in the number of observations, the results are similar.

[Table 10]

We next classify all Arab referees to those exhibiting the most in-group bias and others and classify all Jewish referees to those exhibiting the most in-group bias and others and then examine whether coaches increase the number of Arab starters more when facing the most biased Arab referees relative to when facing the most biased Jewish referees. This entails estimating the following equation:

$$\begin{aligned} Starter_{ijklsr} = \alpha_0 + \alpha_1 ArabPlayer_i * HighBias_r \\ +\alpha_2 ArabPlayer_i * LowBias_r \\ +\theta_{rs}+\beta_{ils} + \gamma_{jls} + \delta_{kls} + \varepsilon_{ijklsr} \end{aligned} \quad (6)$$

where $HighBias_r$ is an indicator variable that takes the value of one when referee *r* is in the top *p* percentiles of the distribution of pro-Arab bias among all Arab referees and the value of 0 otherwise; $LowBias_r$ is an indicator variable that takes the value of one when referee *r* is in the bottom *p* percentiles of the distribution of pro-Arab bias (or equivalently when referee *r* is in the top *p* percentiles of the distribution of pro-Jewish bias) among all Jewish referees; all the other variables are defined as above. We vary the value of the parameter p from 50 in column 2 to 20 in column 5.

The results of the analysis strongly suggest that coaches respond to the anticipated strength of referee in-group bias. In column 2, we find a large difference between the lineup response to Arab referees whose pro-Arab bias is above the median and the lineup response to Jewish referees whose pro-Jewish bias is above the median. The difference is 1.45 percentage points, a value larger than the rule-of-thumb response observed in column 1 of 1.08 percentage points. The size of the response difference increases monotonically as we progressively focus on the Arab and Jewish referees exhibiting the most in-group bias, reaching a level of 2.10 percentage points (or 3.3 percent of the mean) in column 5.

5. CONCLUSION

How do agents who anticipate being affected by bias respond to it? We address this question by studying the decisions of referees and coaches in Israeli football leagues for men. The context offers several key advantages for studying bias. First, the assignment of referees to games is quasi-random. Second, there is a lot of data. Third, the context facilitates the classification of ethnicity not only at the individual level (player and referee) but also, uniquely, at the team level.

Analysis of the data at the player-game level yields robust evidence of in-group bias – the tendency to favor members of one's group – in referee yellow and red card decisions. Estimated bias is not only statistically significant but also economically meaningful.

A well-established finding in the literature on in-group bias is that the strength of this phenomenon tends to increase when group membership is salient. The context we study offers a unique opportunity to study the effect of the salience of ethnicity. Because of the extreme level of residential ethnic segregation in Israel and the fact that teams represent specific localities (or neighborhoods), teams are essentially identified as Arab or Jewish. The context thus naturally generates situations where the salience of ethnicity is particularly high: cross-ethnicity games. Analysis at the player-game level shows that in-group bias in referee penalty card decisions is much stronger in cross-ethnicity games than in same-ethnicity games.

Aggregating the data to the team-game level, we find that referee bias has a meaningful effect on the number of yellow cards, red cards, goals scored, and points awarded to the team. These effects are generally stronger and more statistically significant in cross-ethnicity games.

The results suggest that coaches might want to take into account the ethnic identity of the referee in each game. In particular, bias incentivizes coaches to include in the starting lineups of their teams more players who match the ethnicity of the referee. Analyses of starting lineup decisions indicate that coaches indeed respond to the ethnic identity of the assigned referee in the expected direction.

We show that the strength of the adjustment depends on several factors. First, adjustment is stronger when the team's roster includes sufficiently large numbers of Arab and Jewish players. Second, lineup adjustment is carried out using relatively low-quality players (non-scorers), while high-quality players (scorers) are excluded from it. Third, and probably most important, adjustment is stronger when a stronger bias is anticipated: when the salience of ethnicity is high (i.e., in cross-ethnicity games) and when a more biased referee is assigned to the game.

REFERENCES

Arrow, Kenneth J. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets*, eds. Orley Ashenfelter and Albert Rees, 1-33. Princeton: Princeton University Press.

Bar, Revital, and Asaf Zussman. 2020. "Identity and Bias: Insights from Driving Tests." *Economic Journal*, 130 (625): 1-23.

Bertrand, Marianne, and Esther Duflo. 2017. "Field Experiments on Discrimination." In *Handbook of Field Experiments*, Vol. 1, eds. Abhijit V. Banerjee and Esther Duflo, 309-394. Amsterdam: North-Holland.

Chiappori, Pierre-André, Steven Levitt, and Timothy Groseclose. 2002. "Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer." *American Economic Review*, 92 (4): 1138-1151.

Coate, Stephen, and Glenn C. Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, 83 (5): 1220-1240.

Feld, Jan, Nicolas Salamanca, and Daniel S. Hamermesh. 2016. "Endophilia or Exophobia: Beyond Discrimination." *Economic Journal*, 126 (594): 1503-1527.

Fernando, A. Nilish, and Siddharth Eapen George. Forthcoming. "Peer Pressure and Discrimination: Evidence from International Cricket." *Journal of Law, Economics, and Organization*.

Garicano, Luis, Ignacio Palacios-Huerta, and Canice Prendergast. 2005. "Favoritism Under Social Pressure." *Review of Economics and Statistics*, 87 (2): 208-216.

Glover, Dylan, Amanda Pallais, and William Pariente. 2017. "Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores." *Quarterly Journal of Economics*, 132 (3): 1219-1260.

Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy*, 109 (1): 203-229.

Kudashvili, Nikoloz, and Philipp Lergetporer. 2022. "Minorities' Strategic Response to Discrimination: Experimental Evidence." *Journal of Public Economics*, 208: 104630.

Lang, Kevin, and Ariella Kahn-Lang Spitzer. 2020. "Race Discrimination: An Economic Perspective." *Journal of Economic Perspectives*, 34 (2): 68-89.

Lundberg, Shelly J., and Richard Startz. 1983. "Private Discrimination and Social Intervention in Competitive Labor Market." *American Economic Review*, 73 (3): 340-347.

Mocan, Naci H., and Eric Osborne-Christenson. 2022. "In-Group Favoritism and Peer Effects in Wrongful Acquittals: NBA Referees as Judges." National Bureau of Economic Research Working Paper No. 29860.

Neumark, David. 2018. "Experimental Research on Labor Market Discrimination." *Journal of Economic Literature*, 56 (3): 799-866.

Parsons, Christopher A., Johan Sulaeman, Michael C. Yates, and Daniel S. Hamermesh. 2011. "Strike Three: Discrimination, Incentives, and Evaluation." *American Economic Review*, 101 (4): 1410-1435.

Pope, Devin G., Joseph Price, and Justin Wolfers. 2018. "Awareness Reduces Racial Bias." *Management Science*, 64 (11): 4988-4995.

Price, Joseph, and Justin Wolfers. 2010. "Racial Discrimination among NBA Referees." *Quarterly journal of economics*, 125 (4): 1859-1887.

Shayo, Moses. 2020. "Social Identity and Economic Policy." *Annual Review of Economics*, 12: 355–389.

Shayo, Moses, and Asaf Zussman. 2011. "Judicial Ingroup Bias in the Shadow of Terrorism." *Quarterly Journal of Economics*, 126 (3): 1447-1484.

Welch, Finis. 1967. "Labor-Market Discrimination: An Interpretation of Income Differences in the Rural South." *Journal of Political Economy*, 75 (3): 225-240.

Zussman, Asaf. 2013. "Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars." *Economic Journal*, 123 (572): F433–F468.