Animal Spirits in the Beautiful Game

Testing social pressure in professional football during the COVID-19 lockdown*

Carlos Cueva†

September 11, 2020

Abstract

The COVID-19 pandemic forced almost all professional football matches worlwide to be played in empty stadiums. This large-scale natural experiment offers a unique opportunity to assess the impact of social pressure on decision making and behavior. Using a large dataset from 41 professional football leagues in 30 different countries, I find that the home advantage in match outcomes drops by around one half and that referee bias against away teams disappears following the lockdowns. My results therefore suggest that social pressure exerted by home crowds has an important effect on the behavior of referees and on game outcomes.

1 Introduction

Social pressure can have important effects on behavior according to research in economics, sociology and psychology (e.g. Asch & Guetzkow, 1951; Elster, 1989; Coleman, 1994; Akerlof & Kranton, 2000; Becker & Murphy, 2009; Charness et al., 2007). For example, field experiments in economics suggest that it can affect voting turnout, educational investment, charitable giving, effort in the workplace, consumption and financial investments (for a review, see Bursztyn & Jensen 2017). Theoretically, social pressure and image concerns are thought to play a key role in driving pro-social behavior and shaping social interactions more generally (e.g. Bernheim, 1994; Akerlof & Kranton, 2000; Bénabou & Tirole, 2006). Other theoretical work has emphasized the

*I thank Ignacio Mas Candela for his excellent research assistance and Marianna Battaglia, Lola Collado, Iñigo Iturbe-Ormaetxe, Anna Sanz-de-Galeano and Adam Sanjurjo for valuable comments and suggestions. Financial support from Generalitat Valenciana (SEJI/2019/005) and from the Spanish Ministry of Science and Innovation (Grant PID2019-108193GB-I00) is gratefully acknowledged.

†carlos.cueva@ua.es.
Fundamentos del Análisis Económico (FAE), Universidad de Alicante

importance of social pressure in areas such as labor, education, health, or law enforcement, to name a few (e.g. Kandel & Lazear, 1992; Austen-Smith & Fryer Jr, 2005; Dragone & Savorelli, 2012; Acemoglu & Jackson, 2017).

An important difficulty in testing and interpreting the effects of social pressure on behavior is that there are multiple channels through which it can affect even narrowly self-interested individuals. For instance, yielding to social pressure may bring material benefits in the future, given that most social interactions are repeated. On the other hand, it can provide information about the expected appropriate behavior in unfamiliar environments, or serve as a coordination device (Charness et al., 2007). For this reason, researchers have attempted to isolate the varios factors involved by means of experimental manipulations both in the lab and in the field. However, the scale and scope of these experiments are inevitably limited and typically involve inexperienced individuals making small-stakes decisions.

In an effort to overcome these limitations, Garicano et al. (2005) and subsequent authors studied the effect of social pressure using naturally occurring data from professional sport. The main advantages of professional sport are that (i) it involves highly experienced players making decisions in their field of expertise, (ii) it is a naturally controlled environment, and (iii) there are high stakes and no uncertainty about the payoff structure.[1] These studies find that referees tend to favor the home team in their decisions, and that the degree of favoritism increases with match attendance and decreases in stadiums where the audience is further away from the field (for a review, see Dohmen & Sauermann 2016). Together, the evidence suggests that the reason why referees favor the home team is the social pressure exerted by home crowds. However, since most of this evidence is based on correlations, it is difficult to find conclusive support for the hypothesis that it is the pressure exerted by the audience *per se* which is responsible for this effect.

The COVID-19 pandemic has brought an unexpected opportunity to test this hypothesis. As part of the social distancing measures put in place by governments to limit the spread of the virus, people were barred from attending sporting events for an extended period of time. Association football (soccer), by far the most popular sport in the world in terms of fans and viewership, was to be played in empty stadiums since March 2020 almost everywhere in the globe. This constitutes one of the biggest experiments in sport history and offers a unique opportunity to investigate the causal effect of social pressure exerted by home crowds on a much larger scale than ever before.

My dataset comprises over 230,000 matches from the professional football leagues of 30 countries between 1993 and 2020. 2,749 of these matches are played under full or partial lockdown in a total of 28 countries. I examine the impact of the lockdowns on match outcomes (i.e. the proportion of home wins, draws and away wins) and, referee calls, in particular, the number of fouls, yellow cards and red cards.[2] Before

[1] The use of sports data has been very successful in other areas in economics such as game theory (Chiappori et al., 2002; Palacios-Huerta, 2003; Walker & Wooders, 2001), behavioral economics (Camerer, 1989; Chen et al., 2016; Gilovich et al., 1985; Miller & Sanjurjo, 2018, 2019) and labor economics (Parsons et al., 2011; Szymanski, 2000; Kahn, 2000).

[2] A yellow card (also known as a "booking") is shown by the referee to indicate to a player who has committed a serious offence that he or she has been officially cautioned. A player who receives a

the lockdowns, I find that teams were on average 16 percentage points more likely to win at home than away. After the lockdowns, the gap decreased to 8 percentage points. In other words, the home advantage dropped by around one half. The effect of the lockdowns is even more dramatic when it comes to referee calls. While referees made consistently more calls against away teams than against home teams before the lockdowns, this gap completely disappears after the lockdowns.

Natural experiments in sport, whereby a match has to be played behind closed doors due to an exogenous event, offer a valuable opportunity to investigate the causal effects of home crowds on players and referees. However, these events are rare. Moore & Brylinsky (1993) analyzed 11 college basketball games played behind closed doors due to a measles epidemic, while Pettersson-Lidbom & Priks (2010) and later Reade et al. (2020) analyzed 21 and 160 closed-door matches in professional football, respectively. Overall, these studies find mixed evidence of an effect of home crowds on team performance. While Reade et al. (2020) find partial evidence of a reduction in referee bias, stadium closures in these cases were mostly the result of one-off punishments to home teams for corruption or bad behavior of the audience. One concern is that referees may have been negatively predisposed towards home teams playing in these circumstances, given their recent bad behavior. Another concern is that these teams, or their home fans, may have been engaging in illegal practices that were detected and stopped at the time of the sanctions which were previously helping them win games or sway referees. Both of these mechanisms could potentially lead to an overestimation of the effect of social pressure exerted by the audience.[3]

My findings are in line with Reade et al. (2020), in that home crowds have a particularly strong effect on the way referees treat away teams. However, while their study only finds a significant reduction in the number of referee bookings made against away teams, my results show a large and significant effect of the lockdowns on the *difference* in the number of referee calls made against the home team and the away team, which constitutes a stronger indicator of referee bias. To the best of my knowledge, my results are also the first to provide consistent evidence of a causal effect of home crowds on match outcomes. This result is important because what home fans ultimately care about is whether their team wins or not. Stadium closures are associated with a drop of around 4 percentage points in the winning probability of the home team. Even though this is a sizeable drop, earlier studies analyzing data on stadium lockdowns were underpowered to detect effects of this magnitude.

On a more general level, my results demonstrate that social pressure can be a powerful force able to influence even highly skilled and experienced individuals making

second yellow card in a match is immediately expelled from the game and may not be replaced by a substitute. A red card is either shown immediately after a second yellow card or after a particularly serious offence to indicate that the player has been expelled from the game. Data on referee calls is available only for a subset of this data (88,025 matches). Another popular measure in studies of referee bias introduced by Garicano et al. (2005) is stoppage time. Towards the end of a match, the referee may allocate a small amount of extra time (typically less than five minutes) to compensate for time wasted during the game. Unfortunately this measure is not available in my dataset.

[3]Reade et al. (2020) also includes Italian matches from early 2020 played behind closed doors due to COVID-19 just before the global lockdowns. However, they pool these observations with earlier data on stadium closures to estimate the effect of home crowds on referee bias.

decisions in their field of expertise and, in the case of referees, with strong incentives to remain neutral. This raises questions about the ability of policymakers to overcome the undesirable consequences of social pressure merely by training and incentivizing decision makers appropriately.

2 Data

An observation in this dataset consists of a specific match between two football teams. The team playing in their home ground is called the *home team* and the visitor team is called the *away team*. In football leagues, all teams play against each other twice, once at home and once away. There are three possible relevant outcomes in a football match: a home win, a draw or an away win.

The data used comes from two freely and publicly available sources online (football-data.co.uk and fbref.com) and covers the period 1993-2020. It comprises a total of 233,666 matches from 1,708 football teams playing in 41 different professional leagues from 30 countries. In addition to the final match outcome, a subset of this data from 21 European leagues plus Turkey, comprising 88,025 matches also include the number of fouls, yellow cards and red cards awarded to either team. There are a total of 2,749 matches from 28 countries played under lockdown following the onset of COVID-19. 10 countries in the dataset allowed limited numbers of fans to attend stadiums during part of the lockdowns. Thus, the dataset contains 2,136 matches played under strict lockdown with no fans and 613 matches played under partial lockdown with limited numbers of fans. Table 1 summarizes the data. A timeline of the lockdowns in each country can be found in the online Appendix.

3 Results

I begin by reporting simple aggregate statistics. For this part of the analysis I focus on the period 2012-2020 and exclude leagues with no matches after the lockdowns. I choose this time interval because it provides the broadest league coverage, while keeping the sample of leagues roughly constant across time. Panel A in Table 2 reports match outcome frequencies. Overall, home teams won 45% and 41% of the matches before and after the lockdowns, respectively. Away teams won 29% and 33% of matches before and after the lockdowns, respectively, while 26% of matches ended in a draw both before and after the lockdowns. The difference in the proportion of home wins, away wins and draws before and after the lockdowns is highly significant according to Fisher's exact test ($p < 0.001$).

Turning to referee calls, Panel B in Table 2 displays mean number of fouls, yellow cards and red cards to home and away teams before and after the lockdowns. Among the three, yellow card bookings are perhaps the most interesting disciplinary action to examine. Yellow cards are given for fouls that are serious, but not serious enough to warrant a red card, which entails a fair amount of subjective judgement by the referee. Importantly, yellow cards have a larger negative impact on the booked team than fouls, and therefore tend to provoke a stronger reaction from the home crowd. Even though

Table 1: Matches per football league in the dataset before and after COVID-19 lockdowns.

	Pre-lockdown		Post-lockdown		
Country & League	1993/94–2011/12	2012/13–2019/20	2019/21	Total	F/Y/R
England Football League Championship	10,488	4,308	108	14,904	11,039
Spain Segunda División	7,310	3,574	122	11,006	1,365
England Premier League	7,384	2,948	92	10,424	7,600
Spain La Liga	7,384	2,930	110	10,424	5,700
Italy Serie B	6,528	3,381	113	10,022	1,183
France Ligue 1*	6,850	2,942	16	9,808	5,614
Italy Serie A	6,406	2,910	130	9,446	5,699
France Ligue 2*	6,254	2,930	20	9,204	1,060
Germany 2.Bundesliga	5,888	2,367	81	8,336	1,834
Germany Bundesliga	5,814	2,365	83	8,262	5,814
England National League	3,728	4,309	6	8,043	8,043
Turkey Süper Lig	5,474	2,376	72	7,922	918
Portugal Primeira Liga	5,112	2,226	90	7,428	918
Belgium Jupiler Pro League	4,948	1,904	36	6,888	748
Greece Super League	4,658	1,992	58	6,708	717
Scotland Premiership	3,816	1,775	33	5,624	4,544
Brazil Serie A	0	3,039	74	3,113	0
USA Major League Soccer	0	2,974	109	3,083	0
Mexico Liga MX	0	2,610	74	2,684	0
Japan J1 League*	0	2,470	115	2,585	0
Poland Ekstraklasa*	0	2,224	104	2,328	0
Romania Liga I	0	2,196	83	2,279	0
Sweden Allsvenskan	0	1,936	144	2,080	0
Norway Eliteserien	0	1,935	129	2,064	0
Russia Premier*	0	1,884	112	1,996	0
Denmark Superliga*	0	1,720	76	1,796	0
Finland Veikkausliiga*	0	1,563	74	1,637	0
Ireland Premier Division	0	1,582	26	1,608	0
China Super League	0	1,440	72	1,512	0
South Korea K League 1*†	0	1,368	112	1,480	0
Austria Tipico Bundesliga	0	1,409	63	1,472	0
Switzerland Super League*	0	1,377	67	1,444	0
Bulgaria First League*†	0	1,255	79	1,334	0
Hungary NB I*†	0	1,180	66	1,246	0
England League One	10,488	4,264	0	14,752	10,888
England League Two	10,308	4,304	0	14,612	10,927
Netherlands Eredivisie	5,814	2,374	0	8,188	844
Scotland Championship	3,240	1,397	0	4,637	854
Scotland League One	2,700	1,399	0	4,099	859
Scotland League Two	2,700	1,397	0	4,097	857
Argentina Superliga	0	3,091	0	3,091	0
Total	133,292	97,625	2,749	233,666	88,025

*Countries allowing limited numbers of fans during part of the lockdowns
†Data available from 2014

The column headed F/Y/R displays the number of matches that include data on fouls, yellow cards and red cards. Leagues with no matches after the lockdowns are placed at the bottom of the table. Otherwise, data sorted by the total number of matches per league.

a red card is the most impactful call a referee can make, this happens rarely, which makes a treatment effect harder to detect.

As the table shows, referee calls tend to become significantly less favorable towards the home team after the lockdowns. Home teams tend to concede 10% more fouls, receive 22% more yellow cards and 33% more red cards on average per match. Inter-

Table 2: Match outcomes (Panel A) and mean number of referee calls per match (Panel B) before and after lockdown.

A) Match outcomes

	Pre	Post
Home win	44.87%	41.07%
Draw	26.33%	26.01%
Away win	28.80%	32.92%
N	79,399	2,749

Fisher's exact test, $p < 0.001$

B) Referee calls

	Pre	Post	$p-$value*
Home fouls	12.73	14.08	<0.001
Away fouls	13.16	13.85	<0.001
Home yellows	1.79	2.18	<0.001
Away yellows	2.10	2.09	<0.001
Home reds	0.09	0.12	<0.001
Away reds	0.12	0.12	0.772
N	28,277–31,716	1,164–1,170	

*Mann-Whitney $U-$tests

estingly, referees also call more fouls against away teams after lockdown, although the increase is smaller than for home teams (around 5% more fouls per match). However, when it comes to yellow cards, referees again appear to behave less favorably towards the home team, making significantly fewer bookings against the away team.

Given that pre-lockdown data includes matches since 2012, it is important to account for potentially relevant changes in professional football over this period, notably the introduction of video assistant referee (VAR) between 2017 and 2019. It is possible that VAR has resulted in a reduction in home advantage in recent years, by helping referees make more accurate (and perhaps less biased) decisions. In order to assess whether the above findings are indeed the result of a change in trend triggered by the lockdown and not due to some other recent changes such as the introduction of VAR, I now examine match outcomes and referee calls season by season.

Panel (A) of Figure 1 displays the proportion of home wins, draws and away wins in each season and is based on around 10,000 matches per season.[4] While the proportion of home and away victories is relatively stable during the pre-lockdown period at around 45% and 29%, respectively, a clear jump of around 4 percentage points occurs after lockdown. Another way to look at the size of the effect is that the gap in home and away victories goes down from 16 percentage points to 8 percentage points –in other words, the home advantage drops by around a half. The proportion of post-lockdown match outcomes is significantly different than that of each previous season (Fisher's exact tests, $p \leq 0.005$).

Panels (B–D) of Figure 1 display mean number of fouls, yellow cards and red cards per match. A similar pattern emerges in each of the three figures, which in this case are based on around 3,000–5,000 matches per season. Every season pre-lockdown, home teams receive significantly fewer bookings than away teams, but the difference disappears and even reverses post-lockdown.

In terms of fouls (Panel B), away teams were penalized around 0.4 times (or 3%) more than home teams per match before the lockdowns, and the difference is significant at the 1% level in every season (Wilcoxon signed-rank tests). Post-lockdown, home

[4]Even though the last year in the data is 2020, it is labelled 2019 because these matches correspond to the 2019 season for most leagues in our sample.

Figure 1: Match outcomes and referee calls by season.

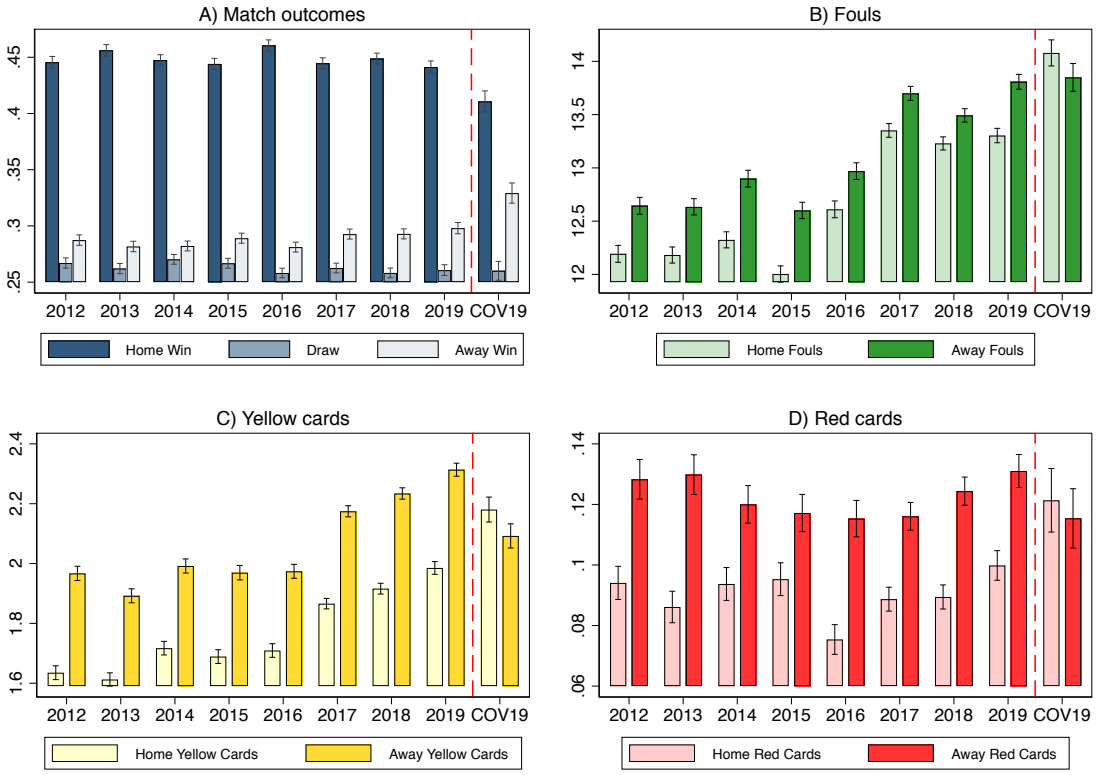

Data from season 2012/13 to season 2020/21. Seasons 2019/20 and 2020/21 are pooled together and split by pre and post-lockdown. Error bars represent ± 1 standard error. Panel A: proportion of home wins, draws and away wins. Panel B: mean number of fouls per match. Panel C: mean number of yellow cards per match. Panel D: mean number of red cards per match.

teams are penalized 0.23 times (or 1.7%) more than away teams and the difference is no longer significant ($p = 0.112$). Comparing differences in home and away fouls post-lockdown with the same measure in each previous season yields a significant difference-in-differences in every pairwise comparison (Mann-Whitney $U-$tests, $p < 0.01$).

Panel C of Figure 1 shows yellow card bookings and highlights an even starker pattern than that of fouls. Before the lockdown, away teams receive around 17% more yellow cards than home teams, a difference which is highly significant in every season (Wilcoxon signed-rank tests, $p < 0.001$). By contrast, home teams received 4% more yellow cards than away teams post-lockdown, although the difference is not statistically significant ($p = 0.113$). Finally, the difference-in-differences between the number of yellow cards pre and post-lockdown is highly significant in all pairwise comparisons (Mann-Whitney $U-$tests, $p < 0.001$).

Panel D of Figure 1 displays the average number of red cards per match and season. The data shows a very similar pattern to yellow cards, although their relative

infrequency make standard errors larger. Away teams receive 36% more red cards than home teams before the lockdowns and the difference is significant in every season (Wilcoxon signed-rank tests, $p < 0.05$), while the gap disappears after the lockdowns ($p = 0.913$). The difference-in-differences in red cards is significant at the 5% level every season except 2014 and 2015 (the former at the 10% level, Mann-Whitney $U-$tests).

3.1 Regression analysis

In this section I check whether the previous results are robust to the inclusion of several controls. As mentioned earlier, some countries started allowing limited numbers of fans back into stadiums after the lockdowns. To control for this, I include both a *lockdown* dummy and a *partial* dummy in my regressions. The lockdown dummy takes value 1 for all matches played after the lockdowns while the *partial* dummy takes value 1 if the match is played after the lockdowns but some fans are allowed to attend. If the effect of the lockdowns on match outcomes and on referee calls is due to the absence of fans, then the estimated effect of the *partial* dummy should partly offset that of the lockdown dummy.

I estimate four different specifications which include various combinations of fixed effects, including league, season, season trimester, team and decade-team interactions. Season trimester fixed effects are included in order to separate the effect of the lockdowns from potential calendar effects (for instance, the lockdowns occurred in the last trimester of the season for most European leagues). Team fixed effects are included separately for home and away teams. Finally, decade-team interactions are the result of interacting decade and team fixed effects, where decades are 1993-2002, 2003-2012 and 2012-2020. This allows me to capture changes in the quality of home and away teams from decade to decade.[5]

Match outcomes are estimated using a linear regression where the dependent variable takes value 1 if the home team wins, 0.5 if there is a draw, and 0 if the away team wins[6]. Estimated coefficients may therefore be interpreted intuitively as changes in the probability of a home win. Table 3 displays the estimated effect of the lockdowns according to four linear regressions in which standard errors are corrected for league-level clusters. The first column of Table 3 shows the estimates of a regression that includes league and season fixed effects. The second column adds trimester fixed effects. The third column is similar to the second, but league fixed effects are replaced by home team and away team fixed effects. Finally, the fourth column is similar to the third, except that it now includes decade-team fixed effects instead of team fixed effects. The estimated effect of the lockdowns according to these regressions is between 4.3 and

[5]Some leagues in our data started season 2020 before August. Since up to now almost all matches from season 2020 occurred during lockdown, the season 2020 fixed effect would be almost perfectly colinear with the lockdown. For this reason, seasons 2019 and 2020 are pooled into the same season fixed effect in order to be able to identify the effect of lockdown in these cases. Including season-team interactions was impractical in my dataset given the large amount of teams and seasons (around 1,700 home teams, 1,700 away teams and 28 seasons). I estimated this model using data from the seasons 2018-2020 only and obtained similar results (available upon request).

[6]Similar results are obtained using logistic regressions instead, although in this case the dependent variable has to be binarized (= 1 if the home team wins, = 0 otherwise)

Table 3: The effect of the lockdowns on home team victory.

	(1)	(2)	(3)	(4)
Lockdown	-0.043***	-0.040***	-0.036**	-0.035**
	(0.011)	(0.011)	(0.011)	(0.011)
Partial	0.032	0.031	0.041*	0.043*
	(0.018)	(0.018)	(0.017)	(0.017)
League f.e .	Yes	Yes	No	No
Season f.e .	Yes	Yes	Yes	Yes
Trimester f.e.	No	Yes	Yes	Yes
HT f.e + AT f.e.	No	No	Yes	Yes
Decade×HT f.e + Decade×AT f.e	No	No	No	Yes
N	233,666	233,666	233,666	233,666

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Linear regressions on home team victory. Dependent variable = 1 if home team wins, = 0.5 if draw, = 0 if away team wins. Standard errors (in parantheses) adjusted for league clusters.

3.5 percentage points and is significant at the 1% or lower. Given that home teams on average won approximately 59% of matches (counting draws as half a win), this constitutes a 6-7% drop in the winning probability of home teams after lockdown. As hypothesized, the estimated coefficient on the *partial* dummy is positive and of similar magnitude (but opposite sign) to the coefficient on the lockdown dummy, indicating that the effect of the lockdowns disappeared in matches where some fans were allowed to attend.

Having established the effect of lockdowns on final match outcomes, I now turn to referee decisions. A referee may favor the home team either by calling fewer fouls on that team, or by calling more fouls on the away team. Since, a priori, it is not clear whether the home crowd may differentially encourage either form of favoritism, I first focus my analysis on the *difference* between the number of referee calls made against the home and away team in each match. The first three columns in Table 4 display the results. The six remaining columns in Table 4 display the effects of the lockdowns on referee calls made against each team separately. All of these regressions adjust standard errors for league clusters and include the same set of controls as the last specification in Table 3, that is, season, trimester, and decade-team fixed effects.

In terms of the impact of the lockdowns on the *difference* between the number of referee calls made against either team, the gap goes up significantly (i.e. increases for the home team relative to the away team) by around 0.5 fouls, 0.4 yellow cards and 0.03 red cards per match after the lockdowns. These are large effects. Before the lockdowns, away teams received 0.6 more fouls, 0.3 more yellow cards and 0.04 more red cards than home teams and these gaps were always significantly different from zero. After the lockdowns, the gap becomes statistically indistinguishable from zero.

As discussed earlier, yellow cards are arguably the most interesting type of referee call given that they have a larger impact on the game than simple fouls but are not as infrequent as red cards. In line with the results in the previous section, the effect of the lockdowns is clearest on the difference in yellow cards. Here, we also confirm that the *partial* dummy has a significant effect in the opposite direction, indicating that the

Table 4: The effect of the lockdowns on referee calls

	HT calls – AT calls			HT calls			AT calls		
	(Fouls)	(Yellows)	(Reds)	(Fouls)	(Yellows)	(Reds)	(Fouls)	(Yellows)	(Reds)
Lockdown	0.500**	0.369***	0.034*	0.846***	0.133	0.021	0.346	-0.236***	-0.014
	(0.138)	(0.065)	(0.015)	(0.201)	(0.070)	(0.011)	(0.221)	(0.060)	(0.011)
Partial	-0.156	-0.705***	-0.101	0.743*	-0.275	0.039	0.899***	0.431	0.139***
	(0.413)	(0.115)	(0.073)	(0.340)	(0.240)	(0.083)	(0.199)	(0.333)	(0.016)
N	83,631	88,024	88,023	83,631	88,024	88,025	83,631	88,025	88,023

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Linear regressions including season, trimester, home team-decade and away team-decade fixed effects and with standard errors adjusted for league clusters. In the first three columns, the dependent variable is the difference between the number of fouls, yellow cards, and red cards given against the home team and the away team per match. In the middle three columns, the dependent variable is the number of fouls, yellow cards and red cards given against the home team. In the last three columns, the dependent variable is the number of fouls, yellow cards and red cards given against the away team

effect of the lockdowns washes away once some fans are allowed into the stadiums.

Examining referee calls against either team separately reveals that home teams receive 0.8 more fouls and away teams 0.2 fewer yellow cards in empty stadiums. These constitute a 6% increase in home team fouls and a 12% drop in away team yellow cards per match. Given the relative magnitude of these effects and the greater impact that yellow card bookings have on games, this evidence suggests that the lockdowns have a more significant effect in the way referees treat away teams than home teams. This observation is consistent with Reade et al. (2020), who find that the number of yellow cards to home teams are unaffected by lockdowns while the number decreases significantly for away teams.

4 Conclusion

This paper has studied the effect of social pressure on behavior in the realm of professional football. The lockdowns triggered by the onset of the COVID-19 pandemic constituted an exogenous shock forcing virtually all professional football leagues in the world to be played behind closed doors. This provides a unique opportunity to investigate the impact of social pressure exerted by football audiences using a larger dataset than ever before.

In terms of match outcomes, my findings show a sizeable and significant reduction in the home advantage during the lockdown. The effect is even more dramatic in the case of referee calls, where the home advantage completely disappears. The most natural interpretation of these findings is that home crowds are able to influence the behavior of the referee in their team's favor. Home fans do this in order to increase the chances of their team winning, which indeed they do. The large effect of the audience on referee calls suggests that the main channel through which the home crowd impacts teams' performance is through its effect on the referee. However, it is possible that the home crowd might increase the probability of their team winning

in other ways, for instance by increasing their morale or decreasing that of the away team.[7] Similarly, we cannot completely rule out the possibility that home crowds only affect the behavior of players and that referees are simply reflecting this change in an impartial way. However, this seems highly unlikely. Given that home fans typically encourage their team to play aggressively and discourage the away team from doing so, research from social psychology suggests that, if anything, home crowds should induce more aggression from home players than from away players rather than the other way around (e.g. Felson, 1982).

The specific channels through which social pressure impacts referee decisions are difficult to disentangle in the present study. One possibility is that referees yield to the pressure of the home crowd and knowingly make inaccurate calls. This is the kind of behavior documented by the influential conformity experiments by Asch & Guetzkow (1951), whereby a proportion of subjects gave obviously incorrect answers in a simple perceptual task in order to agree with the majority view. Another possibility is that social pressure affects referees' perception of the nature and severity of foul play. Because crowd noise tends to be positively correlated with the severity of fouls during a match, referees may use it as a cue to guide their decisions. However, since home crowds react much more strongly to fouls committed by the away team than by the home team, this would bias referees against away teams. In support of this hypothesis, Unkelbach & Memmert (2010) experimentally manipulated the volume of the crowd noise using video recordings and found that referees were significantly more likely to rate fouls as deserving a yellow card in the high volume recordings.

One last aspect of the results worth noting is that home teams continue to have a significant advantage over away teams after the lockdowns. The presence of an audience appears to be responsible for around half of the home advantage. The remaining edge of home teams over away teams might be attributed to other factors such as e.g. travel, venue familiarity or pitch conditions (see Carron et al. 2005).

Could the observed effects be explained by factors other than the absence of an audience? One additional factor which changed during the lockdown was that teams were allowed five substitutions per match rather than the usual three (this was done to help them cope with the increased frequency of matches in the final months of the season). This change may have plausibly affected the number of referee calls, since teams could more easily replace players in order to minimize the chances of anyone receiving a red card. Thus, a reasonable hypothesis is that the increase in the number of allowed substitutions would reduce the number of red cards while increasing the number of fouls and yellow cards. This might explain the fact that both home and away fouls increased during the lockdown. However, since this change in the rules affected both home and away teams equally, there is no reason to expect that it would differentially impact one team over the other. A second factor which changed during the lockdown is that football resumed after a temporary interruption of activity. However, similar interruptions occur before the start of every season and I find no evidence that the home advantage is reduced in the first trimester of the season. A third factor is

[7]Interestingly, however, a recent study using NBA free throw data suggests the possibility of a home *disadvantage*, whereby home players may actually "choke" under the pressure of their home crowd (Böheim et al., 2019).

the coronavirus pandemic itself. Health concerns may have affected the behavior of players and referees during the lockdown, but again, there is no reason to expect that these would have a larger negative impact on the home team than on the visiting team. Finally, matches played under partial lockdown shared all three of these factors except that they allowed some fans back into stadiums. The fact that the home advantage returned in these matches lends further support to the interpretation that the absence of fans was the key factor driving the effect of the lockdowns on home advantage.

References

Acemoglu, D., & Jackson, M. O. (2017). Social norms and the enforcement of laws. *Journal of the European Economic Association*, *15*(2), 245–295.

Akerlof, G. A., & Kranton, R. E. (2000). Economics and Identity. *Quarterly Journal of Economics*, *65*(3), 715–753.

Asch, S. E., & Guetzkow, H. (1951). Effects of group pressure upon the modification and distortion of judgments. *Organizational influence processes*, (pp. 295–303).

Austen-Smith, D., & Fryer Jr, R. G. (2005). An economic analysis of "acting white". *The Quarterly Journal of Economics*, *120*(2), 551–583.

Becker, G. S., & Murphy, K. M. (2009). *Social economics: Market behavior in a social environment*. Harvard University Press.

Bénabou, R., & Tirole, J. (2006). Incentives and prosocial behavior. *American economic review*, *96*(5), 1652–1678.

Bernheim, B. D. (1994). A theory of conformity. *Journal of political Economy*, *102*(5), 841–877.

Böheim, R., Grübl, D., & Lackner, M. (2019). Choking under pressure–evidence of the causal effect of audience size on performance. *Journal of Economic Behavior & Organization*, *168*, 76–93.

Bursztyn, L., & Jensen, R. (2017). Social Image and Economic Behavior in the Field: Identifying, Understanding and Shaping Social Pressure. *Annual Review of Economics*, *9*, 131–153.

Camerer, C. F. (1989). Does the basketball market believe in the hot hand? *The American Economic Review*, *79*(5), 1257–1261.

Carron, A. V., Loughhead, T. M., & Bray, S. R. (2005). The home advantage in sport competitions: Courneya and Carron's (1992) conceptual framework a decade later. *Journal of Sports Sciences*, *23*(4), 395–407.

Charness, B. G., Rigotti, L., & Rustichini, A. (2007). Individual Behavior and Group Membership. *American Economic Review*, *97*(4), 1340–1352.

Chen, D., Moskowitz, T. J., & Shue, K. (2016). Decision-Making under the Gambler's Fallacy: Evidence from Asylum Judges, Loan Officers, and Baseball Umpires. *Quarterly Journal of Economics*, (pp. 1181–1242).

Chiappori, P.-A., Levitt, S., & Groseclose, T. (2002). Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, *92*(4), 1138–1151.

Coleman, J. S. (1994). *Foundations of social theory*. Harvard university press.

Dohmen, T., & Sauermann, J. (2016). Referee Bias. *Journal of Economic Surveys*, *30*(4), 679–695.

Dragone, D., & Savorelli, L. (2012). Thinness and obesity: A model of food consumption, health concerns, and social pressure. *Journal of health economics*, *31*(1), 243–256.

Elster, J. (1989). Social norms and economic theory. *Journal of economic perspectives*, *3*(4), 99–117.

Felson, R. B. (1982). Impression management and the escalation of aggression and violence. *Social Psychology Quarterly*, (pp. 245–254).

Garicano, L., Palacios-Huerta, I., & Prendergast, C. (2005). Favoritism under social pressure. *Review of Economics and Statistics*, *87*(2), 208–216.

Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, *17*(3), 295–314.

Kahn, L. M. (2000). The sports business as a labor market laboratory. *Journal of Economic Perspectives*, *14*(3), 75–94.

Kandel, E., & Lazear, E. P. (1992). Peer pressure and partnerships. *Journal of Political Economy*, *100*(4), 801–817.

Miller, J. B., & Sanjurjo, A. (2018). Surprised by the Hot Hand Fallacy? A Truth in the Law of Small Numbers. *Econometrica*, *86*(6), 2019–2047.

Miller, J. B., & Sanjurjo, A. (2019). A cold shower for the hot hand fallacy: Robust evidence that belief in the hot hand is justified. *IGIER Working Paper No. 518, Available at SSRN*, (pp. 295–303).

Moore, J. C., & Brylinsky, J. A. (1993). Spectator effect on team performance in college basketball. *Journal of Sport Behavior*, *16*(2), 77.

Palacios-Huerta, I. (2003). Professionals play Minimax. *The Review of Economic Studies*, *70*(2), 395–415.

Parsons, B. C. A., Sulaeman, J., Yates, M. C., & Hamermesh, D. S. (2011). Strike Three: Discrimination, Incentives, and Evaluation. *American Economic Review*, *101*(June), 1410–1435.

Pettersson-Lidbom, P., & Priks, M. (2010). Behavior under social pressure: Empty Italian stadiums and referee bias. *Economics Letters*, *108*(2), 212–214.

Reade, J. J., Schreyer, D., & Singleton, C. (2020). Echoes: What Happens when Football Is Played behind Closed Doors? *Available at SSRN Electronic Journal*, *3630130*.

Szymanski, S. (2000). A market test for discrimination in the English professional soccer leagues. *Journal of Political Economy*, *108*(3), 590–603.

Unkelbach, C., & Memmert, D. (2010). Crowd noise as a cue in referee decisions contributes to the home advantage. *Journal of Sport and Exercise Psychology*, *32*(4), 483–498.

Walker, M., & Wooders, J. (2001). Minimax play at Wimbledon. *American Economic Review*, *91*(5), 1521–1538.

5 Appendix

Figure 2: Timeline of the lockdowns