Spring 4-27-2022

How does the crowd affect home field advantage? Evidence from COVID affected seasons in the Top 5 European soccer leagues.

Alexander Hopkins
Macalester College, abmhopkins@gmail.com

How does the crowd affect home field advantage? Evidence from COVID affected seasons in the Top 5 European soccer leagues.*

Alexander Hopkins†
Advisor: Professor Gary Krueger‡
Macalester College

April 27, 2022

Abstract

It is well documented that home field advantage is a significant determinant of team success. The specific mechanism of this advantage is difficult to identify. Is it players' superior knowledge of the home field, the convenience of not having to travel, or the cheering fans of the home crowd? Prior to the COVID-19 pandemic, there was no direct way to isolate the crowd's effects on home field advantage. Due to the pandemic, the top five European soccer leagues barred fans from their stadiums. The pandemic created a quasi-natural experiment to study a crowd's effects on the match outcome and refereeing. Using data from Football Reference and FiveThirtyEight from 2015 through the 2021 season, I use a stadium fixed-effects model to better understand crowds' effect on goal differential (a proxy for match outcome). Similarly, I use a three-way fixed effects model on stadium, season, and referee to test the crowd's effect on the referee. With an extra season of data (previous research only used until the 2020 season), I found that fans are worth 0.23 more home goals than away goals across all leagues. I also found that fans were worth 0.46 more yellow cards in favor of the home team, and the other referee outcomes follow this trend. However, in both cases, league-specific results vary. The results support the hypothesis that fans influence match outcome and refereeing, though any conclusion must be tempered with the disclaimer that COVID changed many factors in addition to the absence of fans.

*I would first like to thank my advisor Professor Gary Krueger, and committee Professors Pete Ferderer and Vittorio Addona for their guidance and support throughout this project. I also want to thank Professor Amy Damon and the 2021 Macalester College Economics Honors cohort for their feedback and camaraderie the past 12 months.

†Economics Department, Macalester College
‡Economics Department, Macalester College

Contents

1 Introduction

Home field advantage is simultaneously one of the most widely accepted and least understood phenomena in sports. Edwards and Archambault (1988) found that commentators make more references to the difficulty of beating a home team than to any other factor, including record, talent, injuries, or momentum. Yet, despite this acceptance of home field advantage, we know very little about its mechanisms. Much of the debate about the mechanisms centers around its potential drivers, be that a cheering crowd, a comfy bed, or a familiar pre-match meal. The struggle in the past for researchers has been in isolating these hypothesized drivers. It is not often that teams play in an empty stadium or share a home stadium (such as AC Milan and Inter Milan), which previously was the only way of isolating the crowd or the stadium. This was true until Covid-19.

The Covid-19 pandemic has provided an opportunity unlike any before. It creates a quasi-experiment in which fans are barred from matches while the season takes place as usual. The comfy bed does not change, nor does the familiar pre-match meal. Beginning in March 2020, leagues were paused as the world entered lockdowns. The German League, the Bundesliga, became the first team to resume playing, but they did so behind closed doors. In the months that followed, other elite European soccer leagues began to resume play, also behind closed doors.

Many papers jumped at the opportunity to analyze the half-season of "ghost-games" (Fischer & Haucap, 2020; Reade et al., 2020). These papers provide early looks at how home field advantage changed during the early part of the pandemic. When these researchers wrote their papers, they could not have predicted that the pandemic would follow the world into the next soccer season. Their analysis misses out on over an entire season of "ghost-games," which I will take advantage of in this paper.

This paper uses data from the top 5 European Leagues: England's Premier League, Spain's La Liga, Germany's Bundesliga, Italy's Serie A, and France's Ligue 1 from the 2015 season through the conclusion of the most recent season in 2021. I will use the data in two ways: to study the effects of the crowd on match outcome and the crowd's impact on the referee to make fair decisions.

2 Literature Review

Home field advantage refers to the consistent occurrence of the home team winning more matches than the away team in a season. In a balanced schedule, where each team plays every other team an equal number of times home and away, one would expect home teams to win 50% of the matches played. However, this is not the case; home teams in almost every sport from the 1980s to the 2000s are found to win more than 50% of matches (Clarke & Norman, 1995; Edwards & Archambault, 1988; R. Pollard, 1986; Ponzo & Scoppa, 2018; Reade & Koyama, 2009). The higher the percentage of home wins indicates more notable home field advantage. This phenomenon has made home field advantage into a popular research and contemporary discussion topic among sports fans, primarily because of the uncertainty surrounding its exact cause, and its implications for teams that can figure out how to improve their advantage and mitigate their opponents when traveling.

Researchers have been studying the effects of home field advantage since before the 1980s, with early research providing potential explanations for its existence (Edwards & Archambault, 1988; R. Pollard, 1986). Pollard (1986) found that different North American sports leagues had varying levels of home field advantage, ranging from 53% in baseball to 65% in the now-defunct North American Soccer League. Edwards and Archambault (1988) corroborate Pollard's results, finding similar levels of home field advantage in North American sports. These two papers provide a series of potential causes for home field advantage, namely: "local crowd support," "travel fatigue," "familiarity with local conditions," "referee bias," "special tactics," "psychological effects," "territoriality," as well as other psychological hypotheses for the mechanism of home field advantage(Edwards & Archambault, 1988; R. Pollard, 1986). Researchers agree that there are likely three main components of home field advantage: crowd support (both encouraging the home team and persuading the referee into giving the home team favorable decisions), familiarity with the stadium, and travel fatigue.

In their review of home field advantage literature, Courneya and Carron (1992) conclude that the "what" (the existence) of the home advantage has received the most attention and more attention needs to be put on the "when" and "why" (the mechanism) of home field advantage. In particular, they think focusing on the "why" will be the most beneficial to improving our understanding of home field advantage. Clarke and Norman (1995) take this advice and, in their paper, try to determine the home field advantage effect for each team. A potential limitation of the understanding of home field advantage was that the home teams could win more than 50% of matches over a season. In contrast, an individual team might lose more than 50% of their home matches - the variance in individual team's results was significant. They believed that the "quality of opposition effect overshadowed the home field effect"(Clarke & Norman, 1995). This leads them to add a measure to control for team ability in their regression. Ultimately, they find that most teams had their home field advantage effect vary widely across years. Their paper began to look at some of the "why" of home field advantage, especially as it pertains to individual teams and their varying ability levels. Their paper would be built upon by numerous other papers, including Carmichael and Thomas (2005).

Carmichael and Thomas (2005) focus on the effects of how teams play as a contributing factor of home field advantage, utilizing the first full season of data compiled by the Opta Index. "The Index provides detailed match statistics that, apart from yellow card and red card disciplinary awards, relate to all touches of the ball by each player itemized according to type (Carmichael & Thomas, 2005)." This influx of data gives them the opportunity to analyze how teams might attack and defend differently depending on if they are playing at home or away. This helps answer part of the "why" of home field advantage – if special tactics play a role. Carmichael and Thomas (2005) find that 26.5% of the estimated goal differential (home goals – away goals) is due to the variance in attack and defense variables. These findings suggest that holding all else equal, teams might choose their style of play based simply on whether they are playing home or away (Carmichael & Thomas, 2005). A possible explanation for this finding is that teams are playing into their expectations. A home loss is more unexpected than an away loss, leading them to play more attacking at home but more likely to sit back and defend when playing away.

This still leaves us with two of the most cited reasons for home field advantage, the crowd effect (R. Pollard & Pollard, 2005; Ponzo & Scoppa, 2018) and the referee (Boyko et al., 2007; Buraimo et al., 2010; Nevill et al., 2002), unanswered. Pollard and Pollard (2005) acknowledge the difficulties of isolating the crowd's effect on a match and the difficulties of isolating the mechanism; is it the home or away team that is affected? Without any statistical analysis, they hypothesize that the crowd does not play a major role in determining the outcome of a match. The level of home field advantage is relatively constant across divisions[1] in European Football, whereas the crowd size differs vastly, leading Pollard and Pollard (2005) to conclude that there is a different factor causing home teams to win at a higher rate. However, Ponzo and Scoppa (2018) disagree with these findings. When analyzing same-stadium derbies in Italy (where all that changes is the crowd make-up), they found that the home team scores about 0.45 more goals than the away team and their winning percent increases by 13(Ponzo & Scoppa, 2018).

Beyond affecting player performance, cries from fans for a foul call, for example, may be influencing how the game is officiated. Consequently, during the first two decades of this millennium, the focus of home field advantage's source was more on the interaction between the referee and the crowd (Anders & Rotthoff, 2014; Boyko et al., 2007; Buraimo et al., 2010; Nevill et al., 2002). The main aspects a referee influences in a match are the fouls and cards they award as well as the amount of stoppage time[2]. The underlying logic behind these studies is that the home crowd influences a referee into making decisions in their favor. In a 2002 study, Nevill, Balmer and Williams (2002) prove that referees are, in fact, influenced by crowd noise. When watching incidents with and without noise from the crowd, referees awarded fewer fouls against the home team than watching the replay silently (Nevill et al., 2002). This gives legitimacy to the conclusion that referees impact the outcome of matches in the English Premier League which Boyko, Boyko and Boyko (2007) find. However, in a paper published a year later, Johnston (2008) challenges Boyko and his colleagues' findings. Using more recent data to replicate their findings, Johnston (2008) primarily focuses on the fact that many Premier League teams play in front of near-capacity crowds every weekend. Including these matches shows the potential for added biases. He argues that biases arise in the relationship between good teams and

[1]In soccer terms, divisions are the different levels of competition. For example, in England, the divisions start in the Premier League and are followed by the Championship, League 1, League 2, National League, etc.

[2]Stoppage time refers to the time that is added on after each half. This time is added for injuries and other match-related stoppages.

their ability to attract large crowds. If a good team is playing in front of a large crowd, is their performance a result of the crowd or their exogenous ability? After condensing the dataset to grounds where there was considerable variation in the data, he found that the referee effect was negligible (Johnston, 2008).

This finding does not mean that we throw out the impact of referees on home field advantage. A more recent study by Anders and Rotthoff (2014) finds that referees in more hostile environments, like the Bundesliga, where a referee may feel a threat to their safety because of historical fan violence, are more prone to making biased decisions than referees in friendlier environments like the MLS.

With Covid interrupted seasons, the studies of crowds and home field advantage exploded (Benz & Lopez, 2021; Davis & Krieger, 2021; Fischer & Haucap, 2020; Hegarty, 2021; Reade et al., 2020). Some of these papers found mixed results, with the change in home field advantage depending on the league (Benz & Lopez, 2021; Fischer & Haucap, 2020), while others reported a drop in home field advantage(Hill & Van Yperen, 2021; Reade et al., 2020). The methods employed and the number of leagues analyzed vary across papers. Some papers use an OLS while others use a bivariate Poisson model while others still choose to study the leagues in one country or the top leagues in multiple. Despite these differences, the models take a similar shape in all the papers: the outcome variable Y is tested before Covid and after Covid while controlling for some mix of team-specific variables. Choices of the outcome variable include home goals, goal differential, points (3/1/0), home yellow cards, yellow card differential, and other in-game actions such as corner kicks and fouls (Benz & Lopez, 2021). For example, Fischer and Haucap (2020) control for the difference in player value, table position (league standing), points earned in the last three matches, travel distance, stadium altitude, within week matches, and matches after 6:00 PM.

These papers provide a comprehensive report of how the absence of crowds has impacted home field advantage. One notable gap in these papers is their exclusion of the 2020-2021 seasons, which were still played almost entirely behind closed doors (save for a few weekends where some stadiums let in a small subset of fans). This extra season of data opens the possibility of more questions. Was the lack of fans an initial shock that normalized as players adjusted to the new normal? The Covid-19 pandemic has given us a once-in-a-lifetime opportunity to study the effects of the crowd on home field advantage as well as the impact of the crowd on the referee in the top European football leagues. My paper will help fill this gap by answering what happened in the entire fan embargo, rather than just the culmination of the 2019-2020 seasons.

3 Theory

3.1 Player Theory

Psychologists have long studied the effects of how spectators impact performance. The first such investigation took place in 1898 when Norman Triplett found that cyclists had faster times competing with other cyclists than when they were racing by themselves (Triplett, 1898). On the other hand, a subsequent study by Pessin (1933) found that students performed worse when performing arbitrary memory tasks when they were observed than when they were not (Bond & Titus, 1983). These two papers formed the backbone for future social facilitation papers as researchers tried to recreate and isolate the causes of Triplett and Pessin's findings.

An integral theory is drive theory, proposed by Robert Zajonc (1965). Zajonc uses terminology from the Hull Spence motivation model to explain the social facilitation effect, particularly as it pertains to familiar versus unfamiliar tasks. The Hull Spence model represents behavior as a function of habit strength and generalized drive (Bond & Titus, 1983). Habit strength refers to the previous conditioning with a behavior, while generalized drive is the amount of energy you have to complete your habits. In their theory, a stimulus evokes a habit, either an individual habit or a set of mutually exclusive habits. Zajonc theorizes that if the stimuli invoke a set of mutually exclusive habits, the generalized drive should multiply by differences in habit strength. The resulting product increases the probability of the dominant response and decreases the likelihood of the subordinate response. The dominant response is often the correct decision for a skilled performer, while for an unskilled performer, the dominant response is the incorrect decision (Zajonc, 1965). Zajonc theorized that performing a task in front of people increases an individual's generalized drive level, enhancing the dominant response – improving simple task performance while inhibiting complex task performance (Bond & Titus, 1983). Zajonc (1980) interpreted this increased drive as an evolutionary response, an increase in preparedness for the unexpected actions of others.

From this perspective, the crowd's presence in soccer increases the players' generalized drive, increasing the dominant response. As all these players are playing in their highest respective domestic league, we can reasonably assume that they are incredibly skilled at playing soccer. In most cases, we would expect the dominant response to outweigh the subordinate response for these players, which could be as small as the technique used to control the ball. For these professional athletes, the crowd enhances their performance. Landers and McCullagh (1976) further corroborate this theory. They found that the presence of a crowd facilitated speed and power tasks and that continuous, fine control accuracy tasks were facilitated only if the task was practiced (Meissner, 1994).

While Cottrell, Sekerak, Wack, and Rittle (1968) disagree with the mechanism behind Zajonc's findings, they provide another potential mechanism for the crowd's effect on players. They propose that it is not merely the presence of others that causes an arousal effect but evaluation apprehension for the task performer. Cottrell et al. (1968) believe that the audience is only drive-enhancing if the performer believes that the audience is evaluating their performance (Meissner, 1994). This is precisely what the crowd is doing in a soccer match, audibly evaluating the players' performance. Whether using Zajonc's or Cottrell's theory, the mechanism remains roughly the same – the crowd enhances the drive of the players, which improves their performance.

During the pandemic-imposed ghost games, there were still people watching while there was no crowd. Each team had their complement of substitutes and coaching staff observing them play, which should be enough to trigger some of the drive theory effect. However, McCullagh and Landers (1976) found that experimental participants performed better as audience size increased (McCullagh & Landers, 1976). This would explain the difference between playing in front of 50 people versus 60,000.

In this paper, I argue that the lack of fans in stadiums during the COVID-19 pandemic was an inverse example of the social facilitation effect—the lack of fans negatively affected players, which was more pronounced for the home team. The home team, previously being used to the positive reinforcement from the home crowd, akin to Cottrell et al.'s interpretation of the drive effect, experienced a more significant decrease in performance relative to their performance with the crowd. Despite knowing ahead of time that fans would be absent, players at this level are used to playing in front of packed crowds. They have likely become accustomed to the generalized drive boost the spectators provide, which could cause the social facilitation effect to be more apparent. Players and managers acknowledge that the conditions and atmosphere is different without fans. After a match behind closed doors, Manchester City's goalkeeper, Ederson, was asked how it felt. "The match itself and the whole environment feels different. It is not the same as playing with a packed stadium" (Olley, 2020). Similarly, before the 2021/2022 Premier League season, Arsenal's manager and former player, Mikel Arteta, said the following about fans' difference: "I think it's the energy, the positivity, the belief, I think its protection. When you are at that ground, and I experienced this myself at the Emirates, and you feel that support and energy, that is something driving you forward. Then it is not thinking, it is pure passion and love for what you do" (Arteta, 2021). Both Ederson and Arteta say that fans provide a different atmosphere. Arteta described it as a protective force, which allowed him to perform better, essentially describing the social facilitation effect.

3.2 Referee Theory

The mechanism for referees is likely different. Most decisions they judge are subjective and unique. Working your decision-making process into muscle memory is hard when they lack the repetition required. A referee has to process an incident and compare it to the rule book definition before giving a ruling, all within a relatively small amount of time. In front of a packed crowd, the referee also has to do their best to provide a fair decision, while 60,000 people shout their own unqualified and biased reviews of the same incident.

Despite being highly skilled at their profession, the fans might inhibit the performance of referees in ways that they don't for players.

Evidence from Nevill, Balmer, and Williams' (2002) study found that referees were influenced by crowd noise towards the home team could add another human element - fear. While players might enjoy the home crowd's support for a full 90 minutes, crowds are much more fickle towards referees, forgetting a previous call that went in their favor if they disagree with the most recent one. This hostile environment could cause referees to make their decisions out of fear of adverse reaction from the crowd, which would weigh decisions in favor of the home team.

4 Data Description

To analyze the crowd effects on home field advantage, I will be using data that I collected from FootballReference. I combined this data with FiveThirtyEight's Soccer Power Index rankings to add a comprehensive metric for team rating.

Football Reference

In recent years, publicly available soccer data has dramatically increased. One of these data sources is FootballReference (FBref), a subsidiary of SportsReference, an American company focused on getting sports statistics to the public. One of the appeals of FBref is its partnership with StatsBomb, a global data provider utilized by clubs around the world, who they work with to provide more advanced statistics, such as expected goals. Professional and casual analysts use FBref as a data source due to its reliability and because its HTML tables allow for easy scraping.

FBref houses data for most European soccer leagues and some African, Asian, and South American leagues. This data encompasses individual player, team aggregate, and individual match data. This paper utilizes their match data records.

Each observation is a match played between two teams in one of the Top 5 European Soccer Leagues. The notable variables of this assembled dataset are goals scored, expected goals accrued during the match, venue, referee, season, league, yellow cards accrued, red cards accrued, penalties, and fouls won. While the dataset has games from the season beginning in 2015 until the season that concluded in 2021, FBref only has expected goals data from the 2017 season onwards.

FiveThirtyEight

FiveThirtyEight began publishing club soccer predictions in January 2017 but has a public dataset that extends from the 2016 season until the 2021 season. At the core of their predictions are a metric originally devised for ESPN, Soccer Power Index (SPI). To create a team's SPI, FiveThirtyEight gives each team an offensive and defensive rating. The offensive rating is the number of goals the team would be expected to score against an average team on a neutral field, while the defensive rating is the number of goals they would be expected to concede. SPI is created from these ratings and represents the percentage of points, 3 points for a win, 1 for a draw, and 0 for a loss a team would earn if they played the same average team repeatedly. Each team has an SPI rating between 0 and 100, with better teams having higher ratings. A team's SPI is updated after every match depending on how it performed compared to expectation (if a team was expected to dominate but narrowly won, their SPI could decrease). As of writing, Manchester City has the highest rating, an SPI of 93.7. SPI will be my control for team strength.

In addition to this metric, FiveThirtyEight's dataset also has its own expected goals value. I will use this to fill in the expected goals data for the 2016 season.

Overall, the dataset encompasses 10,855 matches from 2015 to 2021 across five leagues, the English Premier League, Spanish La Liga, German Bundesliga, Italian Serie A, and French Ligue. 2,234 of the matches took place after the COVID-19 restart, and 2,132 of the matches occurred in empty stadiums.

4.1 Data Summary

Beginning with the dependent variables (Table 1), we can see with a very cursory overview that some evidence of home field advantage is present from the 2015 – 2020 seasons. The variable goal difference is calculated by subtracting the number of goals the home team scored by the number of goals scored by the away side. A positive value means that home teams are scoring more than away teams, which I find. Before COVID, the home team, on average, scored 0.36 more goals than the away team compared to 0.16 after. However, the standard deviation of the goal difference variable is relatively large.

Similarly, the Home Points variable, which measures the number of points the home team would earn from their match result (3 for a win, 1 for a draw, and 0 for a loss), also shows evidence of home field advantage, with the home team earning 1.62 points per match on average before COVID and 1.32 points per match after. There is not much variation in this variable, with only three real values of 0, 1, and 3. Considering the small range, the standard deviation is once again large.

When looking at the referee-related dependent variables, such as yellow card difference, a negative value indicates that the away team is being penalized more often than the home team. In the case of yellow card difference, on average, the home team is penalized with 0.28 fewer yellow cards per match than the away team before COVID. After COVID, this number falls to 0.02 fewer yellow cards per match. Like with the other dependent variables, there is a large standard deviation of 1.71 and 1.7 respectively.

Following the methods used by Ponzo and Scoppa (2018), I created a referee decisions difference variable. The referee decision difference variable is an amalgamation of yellow cards, red cards, and penalties, with red cards weighted at three times that of a yellow card and penalties at five times a yellow card. This variable allows me to look at three potential decisions at once. On average, the home team is penalized with 0.63 fewer referee decisions per match than the away team. There is a large amount of variability in the referee decision difference variable, with both the home and the away team being on the good side of the referee. This extensive range helps explain some of the large standard deviation of 3.66 and 3.72 respectively.

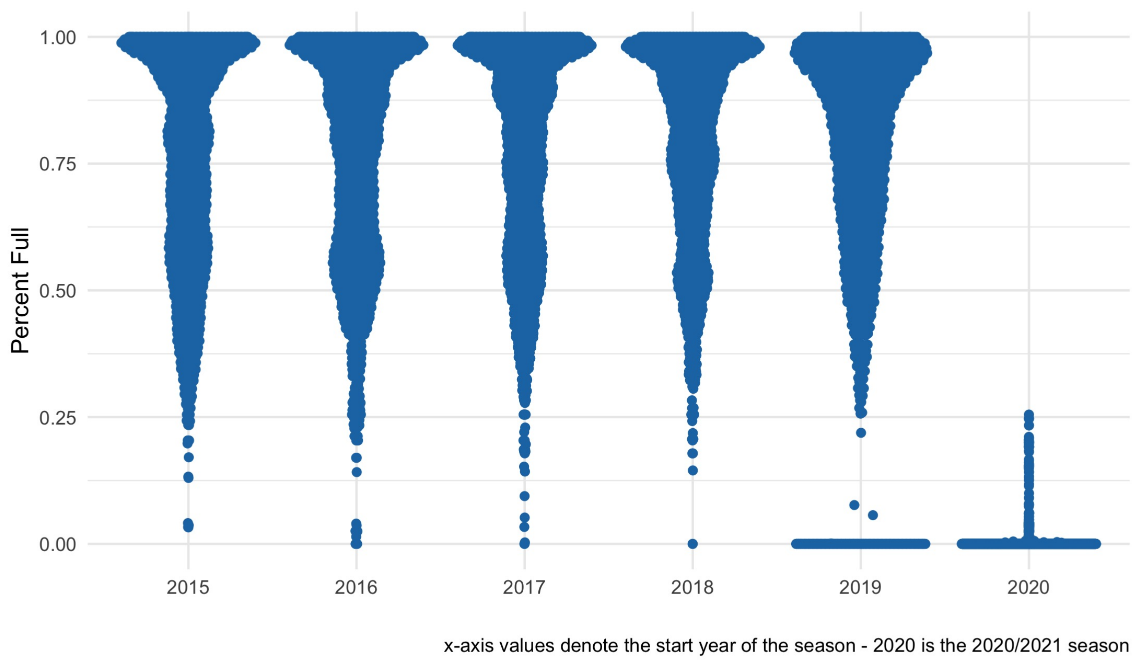

Figure 1: Distribution of Stadium Density

Moving on to the independent variables of interest, in Table 2, all three variables are different ways of measuring fan presence. The first, fans present, is a dummy variable of 1 when fans are present and 0 when they are not. Approximately 20% of the data is from matches with no fans. The second variable, Percent Full, measures stadium density - the attendance divided by the capacity. Originally, the maximum value of of *Percent Full* was greater than 100. This could be due to misreported attendance numbers or a change in the stadium capacity (due to safety or closure for expansion) that is not reflected in the data. These values over 100 were capped at 100.

On average, in matches prior to COVID, teams were filling 75% of their respective stadiums. However, there is a rather large standard deviation of 20%, which indicates a lot of variability in attendance. Some teams might consistently sell out their stadium, while others struggle to fill half of their stadium.

As one might expect, there are considerable variations in stadium capacity across the leagues. Wealthier teams will have larger stadiums, such as Barcelona's Camp Nou, which can seat 99,354. On the other end of the spectrum, some stadiums have a capacity under

Table 1: Descriptive Statistics - Dependent Variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Pre-COVID					Post-COVID				
VARIABLES	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
Goal Difference	8,621	0.361	1.840	-9	9	2,234	0.158	1.847	-7	9
Home Points	8,621	1.617	1.319	0	3	2,234	1.457	1.322	0	3
Yellow Card Difference	8,621	-0.276	1.712	-7	7	2,234	-0.0157	1.697	-6	6
Red Card Difference	8,621	-0.0292	0.444	-3	2	2,234	-0.0152	0.406	-2	2
Referee Decision Difference	8,240	-0.626	3.658	-22	17	2,234	-0.133	3.721	-14	14

Table 2: Descriptive Statistics - Independent Variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Pre-COVID					Post-COVID				
VARIABLES	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
Capacity	8,621	39,225	19,564	7,638	99,354	2,234	38,946	19,071	6,000	99,354
Fans Present	8,621	0.998	0.0456	0	1	2,234	0.0537	0.226	0	1
Percent Full	8,621	76.94	20.73	0	100	2,234	0.469	2.635	0	25.51

Percent Full and Stadium Capcity exclude the 2000 matches that occurred since the COVID pause

Table 3: Descriptive Statistics - Control Variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Pre-COVID					Post-COVID				
VARIABLES	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
Home Soccer Power Index	6,795	67.10	11.65	36.40	96.57	2,234	69.15	11.04	40.49	95.61
Away Soccer Power Index	6,795	67.01	11.65	36.35	96.69	2,234	69.05	11.06	39.36	95.32
Distance	8,621	376.3	285.1	0	2,549	2,234	349.5	223.8	0	1,054
xHRDF	8,241	2.812	0.843	0.842	5.684	2,234	3.048	0.807	0.842	5.579
xHRDA	8,241	3.411	0.919	1.235	7.684	2,234	3.281	0.959	1	6.684
xARDF	8,241	3.410	0.908	1.471	6.947	2,234	3.285	0.879	1.579	6.947
xARDA	8,241	2.812	0.821	1	5.368	2,234	3.048	0.840	1	5.632
Home Rest	8,374	7.626	3.888	2	37	2,129	6.447	3.289	2	23
Away Rest	8,374	7.623	3.894	2	36	2,129	6.449	3.272	2	23

538 only started producing SPI data during the 2016 season

Higher rest days are a result of the Bundesliga's month long mid-season holiday break

xHRDF - Expected Home Referee Decisons For, xHRDA - Expected Home Referee Decisons Against

xARDF - Expected Away Referee Decisons For, xARDA - Expected Away Referee Decisons Away

10,000. The standard deviation of 19,546 reflects this variance.

Looking at some of the potential controls, Table 3, such as Soccer Power Index (SPI), we observe a healthy variation in the ranking of the teams. While there are some outliers on either end of the spectrum, most of the data is clustered between 50-70, which is about the level we expect from teams playing in the top 5 European leagues.

Another important control is the distance traveled by the away team. In a straight line, the average distance traveled by away teams in each match is 370.78km. Teams in some leagues, such as La Liga, must travel further because of the country's size. For example, teams in La Liga must visit Las Palmas in the Canary Islands, which can mean traveling over 2000km. Other teams benefit from sharing a stadium, which results in a travel distance of 0km.

The amount of rest between league matches that each team experiences are consistent. During the normal course of a season, teams usually play on the weekend, which means that the average rest of 7.38 days makes sense. However, these rest days may be extended or reduced during some parts of the season. The Premier League and the Bundesliga take opposite approaches during the holiday period. The Premier League ramps up its fixtures, with teams playing multiple matches in a week, with some teams playing on as little as two days of rest. On the other hand, the Bundesliga pauses its season, giving teams over a month of rest.

Figure 2: Relationship between SPI and points won when playing at home

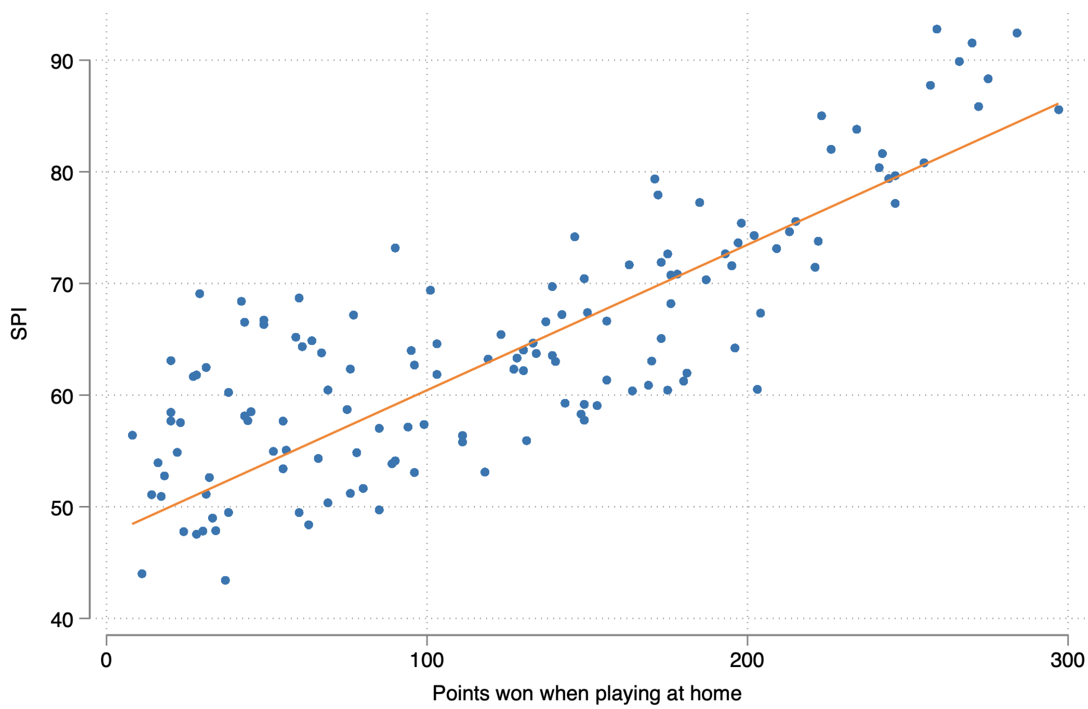

A t-test of goal difference grouped by fans' presence shows an absolute difference between the means of goal difference with and without fans. The significant t-test difference of -0.219 means that we can reject the null hypothesis that the two means are the same.

Table 4: T-Test of Dependent Variables

Variable	N	(1) 0 Mean/SE	N	(2) 1 Mean/SE	T-test Difference (1)-(2)
Goal Difference	2132	0.143	8723	0.362	-0.219***
		(0.040)		(0.020)	
Home Points	2132	1.450	8723	1.617	-0.167***
		(0.029)		(0.014)	
Yellow Card Difference	2132	0.009	8723	-0.279	0.288***
		(0.037)		(0.018)	

Notes: The value displayed for t-tests are the differences in the means across the groups. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

The t-test of home points, when grouped by the presence of fans, shows very similar results. There is no overlap between the two means' confidence intervals, and the t-test is again low enough that we can reject the null hypothesis. These results give credence to the hypothesis that crowds impact home field advantage. The t-test of yellow card difference with fans and without fans, like the two previous variables is significant at the 1% level. The difference in mean yellow card difference with and without fans is significantly different.

5 Empirical Approach

5.1 Match Outcome Model

Following the conceptual model laid out by Clark and Norman (1995) and others, I propose that the outcome of a soccer match can be modeled as a function of home team ability, away team ability, home field advantage, and luck (random error). The home field advantage term is most clearly connected to the social facilitation effect by combining this with the theory above. The social facilitation effect is a positive effect with fans and a negative effect without fans.

$$\text{Match Outcome} = \text{F(home team ability, away team ability, home advantage, luck)}$$

Given this conceptual model and the data available, the actual model follows a similar form albeit with proxies for each of the hypothesized components:

$$Goal\ Difference = \beta_0 + \beta_1(Fans) + \beta_2(SPI_{home}) + \beta_3(SPI_{away}) + \\ \beta_4(Days\ of\ Rest_{home}) + \beta_5(Days\ of\ Rest_{away}) + \beta_6(Distance) + \alpha_i + \delta_t + \epsilon$$

In words, the goal difference (home goals – away goals) of a match can be predicted by the presence of fans, the SPI of the home team, the SPI of the away team, the respective days of rest for the home and the away team, and the distance traveled by the away team. This model is a two-way fixed effects model: α_i, denotes the stadium fixed effects, while δ_t denotes the season fixed effects, ϵ is the error term. A league fixed effect is unnecessary because it would be redundant with the stadium fixed effects. The stadiums are all league specific; Premier League stadiums are not used in any other league and vice versa. The stadium fixed effects will absorb any unobserved differences between the leagues.

As this paper attempts to show the effect of the crowd on the match outcome, I believe it makes sense to use the dummy variable *Fans* which indicates if fans were present at a match. With the large variation in stadium capacities, using raw attendance numbers as a predictor of interest could introduce more biases than solutions. Alternatively, a measure of stadium density, such as *Percent Full* (from Table 1), could also provide a more unbiased, compared to attendence, measure of fan presence.

When the dummy variable is equal to 1, which indicates that fans are present, home field advantage is intact, there is a positive social facilitation effect. However, when there are no fans and the dummy variable is equal to 0, an element of home field advantage is potentially removed leading to a negative social facilitation effect on this term.

538's Soccer Power Index (*SPI*) is my chosen control of team ability. One of the appeals of this metric is that it is a single number, which increases its interpretability. An alternative would be to use two metrics for each team, a measure of attacking and defensive strength. The mean of expected goals accrued by a team over a season could be used to measure attacking strength. *SPI* stands out in this regard, as it provides an accurate team rating at each match in time, whereas the above metric relies on a season long average. *SPI* is relatively sticky over the course of a season; a team's rating might only change by 2 points.

Perhaps the only (somewhat) easily controllable measure of home field advantage is travel fatigue, measured by the number of rest days between matches and travel distance. In league play, rest times are consistent, teams usually play once every weekend. However, some teams might also have midweek fixtures[3] due to their commitments in other European-wide leagues, like the Champions' League or the Europa League. In my dataset, I only have domestic league matches, which means that for some teams, in some games, the rest days are not accurate. To help mitigate this, dummy variables indicate whether a team was competing in, and had not yet been eliminated from, one of these additional leagues when they played their match. One immeasurable issue that arises is play during international breaks. As international competing players disperse to their own countries to

[3]Champions League matches are typically played on Tuesdays and Wednesdays while Europa League matches happen on Thursdays

play other international nations, they play a varying number of matches with differing days of rest between those matches. While their club team might not have played for 14 days, they could have played as recently as three days ago, rendering the original rest variable incorrect for the individual.

The first fixed effect, α_i, this model employs is one on stadiums. In most top-flight domestic soccer leagues, teams play all their home matches at one stadium. But these stadiums are not standard – they vary in capacity, pitch size, altitude and in some cases, have running tracks that separate the pitch from the stands. Stadium fixed effects will control for these differences. They will also control for any time invariant size of the stadium due to closures or expansions.

The second fixed effect, δ_t, is over season. In each season, there is the potential for rule changes within leagues. A good example of this is the introduction of Video Assistant Referee (VAR). VAR was introduced to reduce substantial human refereeing errors from affecting the match outcome. The five leagues of interest all adopted the new technology in different seasons and is operated slightly differently in each league. The season fixed effects will control for changes across seasons, like VAR.

5.2 Referee Outcome Model

The other side of the crowd support question is their impact on the referee. While referee decisions are normally considered individual actions, the model will treat them as an aggregate per match. This relationship can be modeled in a similar fashion to the match outcome equation:

$$\text{Referee Decision} = \text{F}(\text{discipline}_{\text{home}}, \text{discipline}_{\text{away}}, \text{home field advantage, human error})$$

With this conceptual model, the actual model can be created, once again with proxies for the terms above.

$$\begin{gathered} Referee\ Decision\ Variable = \beta_1(Avg\ Referee\ Decision\ Variable\ Received)_{home} + \\ \beta_2(Avg\ Referee\ Decision\ Variable\ Earned)_{home} + \\ \beta_3(Avg\ Referee\ Decision\ Variable\ Received)_{away} + \\ \beta_4(Avg\ Referee\ Decision\ Variable\ Earned)_{away} + \beta_5(Fans) + \alpha_i + \delta_t + \phi_i + \epsilon_{(i,t)} \end{gathered}$$

In words, the difference between the chosen referee variable (e.g., home yellow cards – away yellow cards) of a match can be predicted by the presence of fans and a slightly more complicated set of controls: the average number of referee decisions received (or earned[4])

[4] If the home team earns a yellow card, they have drawn a foul that results in the away team being shown a yellow card.

by the home (or away) team during a match[5]. This model is a three-way fixed effects model: α_i denotes the stadium fixed effects, δ_t denotes the season fixed effects, ϕ_i denotes the referee fixed effects, and ϵ is the error term.

Most of the variables included in this equation have been outlined in the match outcomes model above, such as *Fans* and the two fixed effects, α_i, δ_t, stadium and season. The rationale for their inclusion is the same for the referee. The two new elements are the controls and referee fixed effects.

The team controls take the form of team strength controls that previous researchers have employed (Boyko et al., 2007). For some teams, committing fouls are part of their tactics. For example, Manchester City are known for committing frequent "professional fouls" to disrupt the opposition. The variation in a team's propensity to commit fouls (and, as a result, receive yellow cards) could be large. To control for this, I take the average of a team's transgressions received and earned when they are playing home and away.

Another necessary addition is the introduction of a third fixed effect, ϕ_i, which is a fixed effect on the match referee. As referees are human, they may have different thresholds for what is deemed a foul, yellow card, or red card offense. The referee fixed effect will control for these inherent human differences.

Unlike match outcome, there is no clear variable to choose from for the referee decision. As such, I use four different variables to more holistically observe the impact that fans have on the referee: fouls, yellow cards, red cards, and an aggregate variable r*eferee decisions*. The aggregate variable combines yellow cards, red cards, and penalties into one variable to measure the number of substantial decisions a referee makes each match. Because red cards and penalties are so infrequent, they are weighted three times that of a yellow card. While arbitrary, this variable lets us observe the impact of fans on multiple referee decisions at once.

$$\begin{aligned}\text{Referee Decision Variable} = \text{Yellow Cards Conceded} + 3 * \text{Red Cards Conceded} \\ + 5 * \text{Penalties Conceded}\end{aligned}$$

6 Results

6.1 Fans and Results

Using the two-way fixed effects model to predict goal difference that I outlined in section 5.1, we can test the impact that fans have on the outcome of matches in terms of how many more (or fewer) goals the home team will score compared to the away team. Table 5 shows the results of this regression, first with the leagues combined and then separated.

[5]Each home team receives a pair of these variables for each season, the average number of decisions they received and the average number of decisions they drew. The same is true of each away team. As a result, each team has four variables for each season.

Figure 3: Ratio of Home Points to Away Points in Each League Across Each Season (Mean in red)

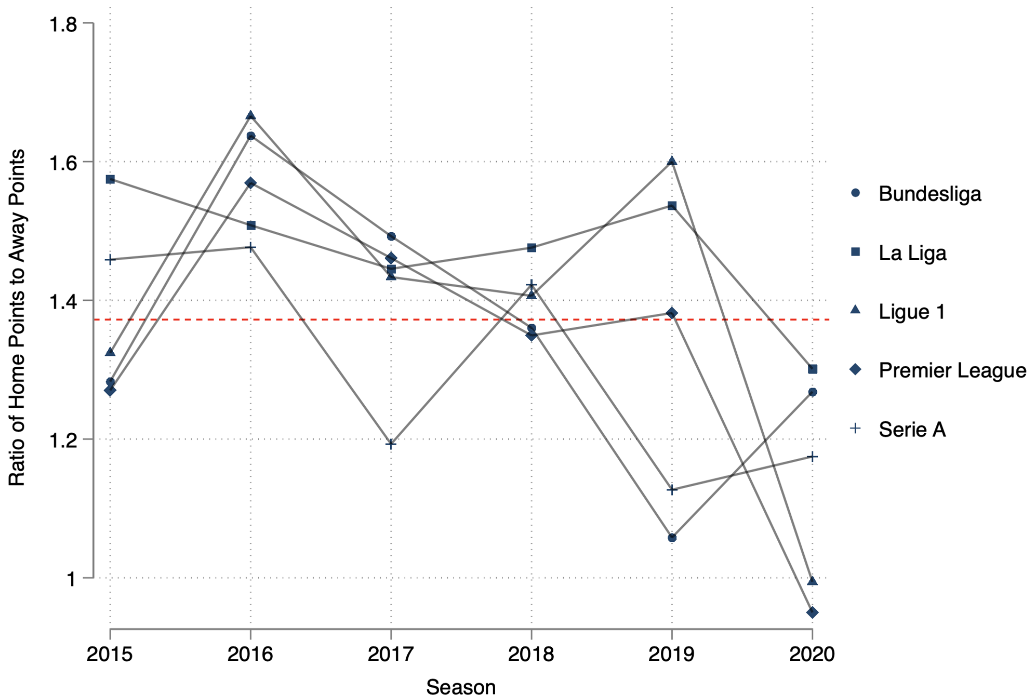

This allows us to see the effects that fans have on matches throughout the top 5 leagues in Europe and any league-specific differences. It is worth remembering that a positive goal difference is in favor of the home team and a negative goal difference is in favor of the away team. A positive coefficient signifies how many more goals the home team scores than the away team, while a negative coefficient signifies how many few goals the home team scores than the away team.

Across the entire dataset, fans were worth 0.174 home goals more than away goals while controlling for team strength, distance traveled, and rest. This effect varies when looking at the leagues individually. The impact of fans is more pronounced in the Bundesliga, La Liga, and Ligue 1, with fans worth 0.265, 0.354, and 0.331 home goals more than away goals, respectively. However, in the Premier League, fans are only worth 0.009 more home goals. On the opposite end of the spectrum, fans in Serie A are detrimental to the home team according to this regression, with their being worth -0.05 more home goals (or 0.05 away goals). However, only the Top 5 Leagues and La Liga results have significant p-values at the 10% threshold.

Looking at the coefficients for the control variables, many make intuitive sense. Across the six regressions, the home and away soccer power index (SPI) values match the expectations. The stronger the home team is (higher home SPI), the expectation is that the goal difference would increase positively, which the coefficients reflect. The opposite is true with away team strength, where the expectation is decreased goal difference. Distance is a control that has mixed effects, likely due to the differences in country size. Regardless of the differences, the effect is minimal. At most, traveling 100 km is worth 0.033 more home goals in La Liga.

The theory behind the inclusion of variables that indicated whether a team was in the Champions League or Europa League was that these teams might have less actual rest in between matches. The coefficients do not necessarily reflect this hypothesis. The expectation would be for the coefficients to appear how they do for the Home Team in the Champions League and the Away Team in Champions League coefficients in the Bundesliga column. If the home team is in the Champions League, the away team has an advantage of 0.307 goals (significant p-value), while if the away team is still playing in the Champions League, the home team has an advantage of 0.079 more goals. While they have insignificant p-values, the coefficients for Ligue 1 in this category are the opposite of the expectations.

The lack of significant p-values on the season fixed effect results led me to rerun the regression without a fixed effect on season. This could be due to the seasons fixed effects variable confounding with the fans dummy variable. This decision was further supported by an F Test that showed an insignificant effect of season on goal difference (F4,141 = 1.16, $p < 0.3327$) and the evidence that the two variables were highly correlated (PW Corr: -0.6664). The overwhelming majority of matches played without fans happen during the 2020 season, meaning that the season fixed effects is not as necessary as I initially expected. To account for the one notable rule change across the seasons (Video Assistant Referee), I introduced a dummy variable that measured whether VAR had been implemented.

$$GoalDifference = \beta_0 + \beta_1(Fans) + \beta_2(SPI_{home}) + \beta_3(SPI_{away}) + \\ \beta_4(Days\ of\ Rest_{home}) + \beta_5(Days\ of\ Rest_{away}) + \beta_6(Distance) + \beta_7(Video\ AR) + \alpha_i + \epsilon$$

Table 6 shows the results from this second regression. The results show much more significance for the predictor of interest, the dummy variable for fans, but signify the same results as above. In the entire dataset, fans are worth 0.212 more home goals and worth as much as 0.463 more home goals in Ligue 1. In La Liga, Ligue 1 and the Premier League fans are worth almost a quarter of a goal to the home team, values significant at the 5% level. However, in the Bundesliga and Serie A, fans are barely worth a 10th of a goal, 0.085 and 0.034 more home goals, respectively. These values also do not have significant p-values.

While it is necessary to fully account for the difference in the amount of rest between matches that teams have, the coefficients on the Champions League and Europa

Table 5: Goal Difference - Season and Stadium Fixed Effects

VARIABLES	(1) Top 5 Leagues	(2) Bundesliga	(3) La Liga	(4) Ligue 1	(5) Premier League	(6) Serie A
Fans Present	0.174*	0.265	0.354**	0.331	0.009	-0.050
	(0.092)	(0.309)	(0.157)	(0.541)	(0.161)	(0.140)
Home Soccer Power Index	0.027***	0.017	0.026**	0.037**	0.019**	0.030**
	(0.006)	(0.015)	(0.011)	(0.017)	(0.009)	(0.012)
Away Soccer Power Index	-0.054***	-0.062***	-0.053***	-0.049***	-0.055***	-0.054***
	(0.002)	(0.007)	(0.005)	(0.006)	(0.004)	(0.005)
Distance in 100KM	0.013	-0.004	0.033**	-0.026	0.004	0.025
	(0.009)	(0.029)	(0.014)	(0.019)	(0.040)	(0.016)
Home Rest	-0.000	0.005	-0.006	0.002	-0.000	0.000
	(0.002)	(0.005)	(0.005)	(0.004)	(0.005)	(0.004)
Away Rest	-0.002	-0.006	-0.000	-0.002	-0.003	-0.001
	(0.002)	(0.006)	(0.005)	(0.004)	(0.006)	(0.003)
Home Team in the Champions League	-0.039	-0.307**	-0.057	0.063	0.039	0.003
	(0.064)	(0.114)	(0.119)	(0.146)	(0.120)	(0.179)
Away Team in the Champions League	-0.141**	0.079	-0.241**	-0.376	-0.013	-0.170
	(0.071)	(0.194)	(0.115)	(0.224)	(0.152)	(0.157)
Home Team in the Europa League	0.038	0.162	-0.084	0.084	0.147	-0.055
	(0.062)	(0.171)	(0.079)	(0.155)	(0.150)	(0.168)
Away Team in the Europa League	0.014	0.028	0.066	-0.245*	0.198	-0.045
	(0.065)	(0.180)	(0.093)	(0.140)	(0.145)	(0.176)
Constant	2.033***	3.476***	1.953**	1.054	2.816***	1.840**
	(0.392)	(1.111)	(0.864)	(1.110)	(0.858)	(0.758)
Observations	9,029	1,530	1,900	1,799	1,900	1,900
R^2	0.244	0.227	0.239	0.225	0.269	0.278
Stadium-FE	YES	YES	YES	YES	YES	YES
Season-FE	YES	YES	YES	YES	YES	YES
r2_a	0.231	0.207	0.221	0.206	0.251	0.263
F	83.74	41.50	29.49	51.11	90.35	51.20
rss	23303	4677	4307	4559	5066	4553

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

League variables do not agree with the expectation because they are closely related to team strength. Teams that participate in the Europa League and Champions League earn the right to based on their performance in the previous season. For example, in the Premier League, the teams that finish in the top 4 qualify for the Champions League, while the 5 and 6 place teams qualify for the Europa League. This could explain why most of the Away Team in Champions League coefficients are negative. The away team playing in the Champions League is a confounder; it is one of the stronger teams in that league, so the coefficients signify more goals for the away team holding all else equal.

The one-way fixed effects model is better for answering the question about fans and their impact on the match outcome, as evidenced by the more significant results below.

Figure 4: Goal Difference Model 1: Coefficient Plot of Dummy Variable for Fans

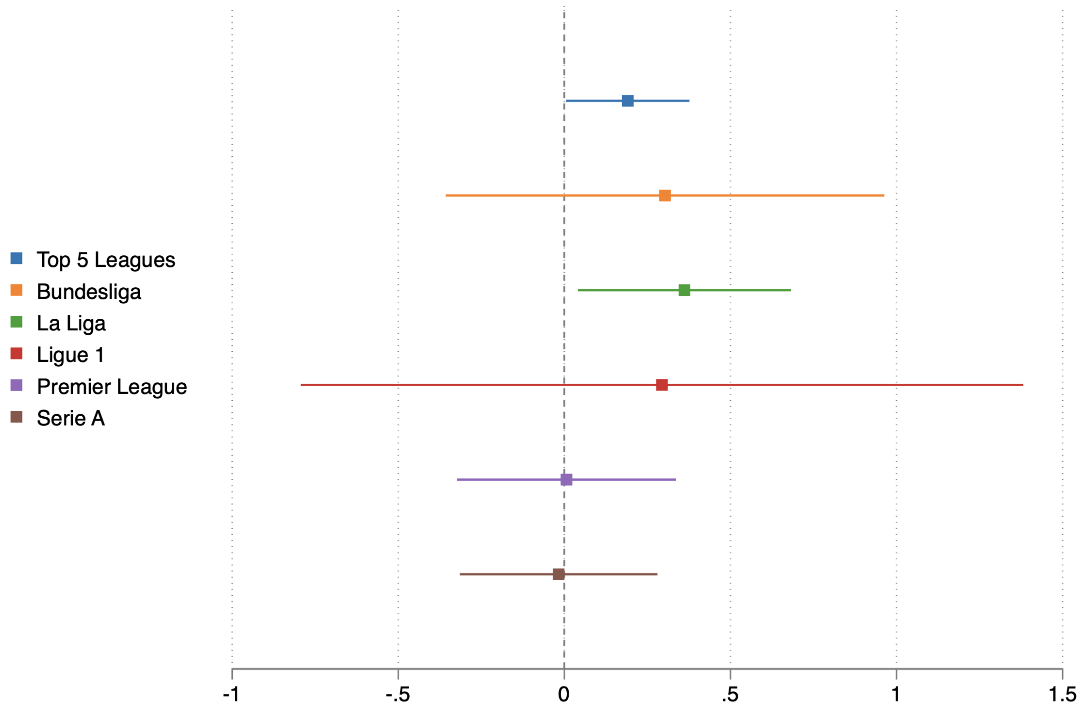

6.2 Fans and the Referee

This paper's second question is about the crowd's impact on the referee's ability to make unbiased decisions. Using the empirical model outlined above, Table 7 shows that referees do make different decisions depending on if fans are present. The values of the coefficients with referee outcomes are the opposite to match outcome dependent variables. A negative coefficient signals that the home team is favored, and a positive coefficient signals that the away team is favored. The referee-related dependent variables are deterrents; teams and players try to avoid collecting them, whereas they are actively trying to score goals. With the current understanding of home field advantage, negative results are expected because the home team is penalized less than the away team, hence a negative difference.

With this interpretation in mind, we can see in Table 7 that referees give 0.427 more yellow cards per match to away teams when fans are present when controlling for both teams' propensity to be awarded yellow cards and their tendency to be punished with them. This coefficient has a p-value that is significant at the 1% level. Almost half of a yellow card more is a practically significant result. A player with a yellow card may be

Figure 5: Goal Difference Model 2: Coefficient Plot of Dummy Variable for Fans

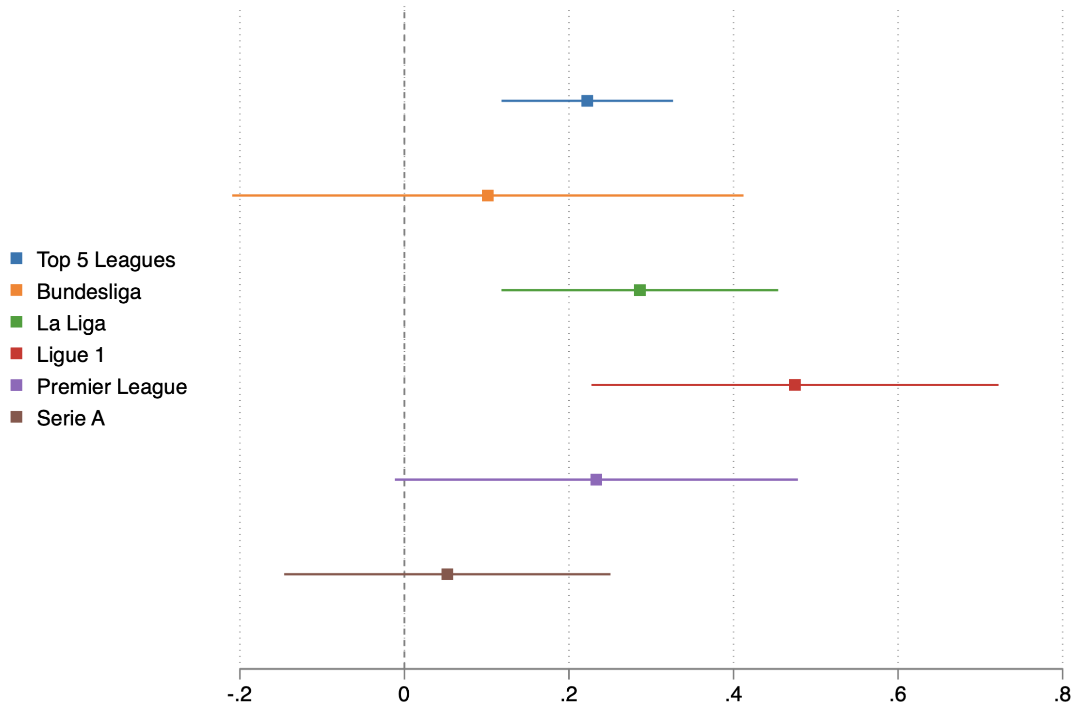

more cautious in their subsequent play in fear that another action might result in another yellow card, which would result in a red card and their team finishing the match with ten players. This number also signifies that the presence of fans has a direct impact on the way the referee manages the match.

In the other three variables tested, it is evident that the same result is true. Referees award more red cards and fouls to the away team than the home team when fans are present, with 0.051 red cards and 0.908 fouls, respectively. While both values once again have significant p-values at the 1% level, the magnitude of the red card difference is negligible. Red cards are a rare occurrence (0.202 red cards per match), so this small magnitude was expected and gives additional evidence to the necessary weighting in the referee decision variable. With an average of more than 25 combined fouls in a match, less than one additional foul is a small relative magnitude. However, in a low-scoring game like soccer, one foul could significantly affect the match outcome. These results further show that fans influence the referee.

Table 6: Goal Difference - Stadium Fixed Effects

VARIABLES	(1) Top 5 Leagues	(2) Bundesliga	(3) La Liga	(4) Ligue 1	(5) Premier League	(6) Serie A
Fans Present	0.212***	0.086	0.273***	0.463***	0.242*	0.034
	(0.052)	(0.149)	(0.081)	(0.121)	(0.121)	(0.092)
Home Soccer Power Index	0.029***	0.016	0.025**	0.037**	0.024***	0.032***
	(0.005)	(0.014)	(0.011)	(0.017)	(0.008)	(0.010)
Away Soccer Power Index	-0.054***	-0.062***	-0.053***	-0.049***	-0.054***	-0.053***
	(0.002)	(0.007)	(0.005)	(0.006)	(0.003)	(0.005)
Distance in 100KM	0.012	-0.005	0.033**	-0.025	0.005	0.024
	(0.009)	(0.029)	(0.015)	(0.019)	(0.041)	(0.015)
Home Rest	0.000	0.005	-0.006	0.002	-0.000	0.000
	(0.002)	(0.005)	(0.005)	(0.004)	(0.005)	(0.004)
Away Rest	-0.002	-0.006	-0.000	-0.002	-0.003	-0.001
	(0.002)	(0.006)	(0.005)	(0.004)	(0.006)	(0.003)
Home Team in the Champions League	-0.037	-0.331**	-0.048	0.081	0.062	-0.005
	(0.064)	(0.119)	(0.120)	(0.143)	(0.129)	(0.180)
Away Team in the Champions League	-0.145**	0.063	-0.231**	-0.357	-0.026	-0.182
	(0.070)	(0.195)	(0.108)	(0.218)	(0.144)	(0.154)
Home Team in the Europa League	0.036	0.128	-0.083	0.069	0.166	-0.062
	(0.063)	(0.171)	(0.076)	(0.161)	(0.147)	(0.160)
Away Team in the Europa League	0.012	-0.007	0.070	-0.254*	0.195	-0.058
	(0.064)	(0.175)	(0.093)	(0.144)	(0.141)	(0.177)
VAR Implemented	-0.033	-0.228	0.057	-0.001	0.035	-0.065
	(0.055)	(0.134)	(0.063)	(0.161)	(0.105)	(0.067)
Constant	1.828***	3.659***	2.079**	0.856	2.169***	1.574**
	(0.380)	(1.076)	(0.883)	(1.009)	(0.691)	(0.726)
Observations	9,029	1,530	1,900	1,799	1,900	1,900
R^2	0.244	0.226	0.239	0.224	0.266	0.277
Stadium-FE	YES	YES	YES	YES	YES	YES
Season-FE	NO	NO	NO	NO	NO	NO
r2_a	0.231	0.207	0.222	0.207	0.250	0.263
F	105.9	25.80	36.81	41.33	83.46	40.93
rss	23313	4687	4309	4564	5082	4559

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fans are worth 0.602 more referee decisions[6] in favor of the home team, Table 1 and has a significant p-value at the 1% level. In matches without fans, referees gave more than half a decision fewer to the home team than they did to the away team.

6.2.1 Yellow Cards

As with the goal difference analysis above, the composite results show a significant relationship between fans and the difference in yellow cards, but the league-specific effects

[6] Referee Decision = Yellow Cards Conceded $+ 3 *$ Red Cards Conceded $+ 5 *$ Penalties Conceded

Table 7: Referee Decisions Differences - Top 5 Leagues

VARIABLES	(1) Yellow Card Difference	(2) Red Card Difference	(3) Foul Difference	(4) Referee Decision Difference
Fans Present	-0.427***	-0.051***	-0.908***	-0.602***
	(0.061)	(0.018)	(0.213)	(0.174)
Constant	0.318	0.121***	1.788**	0.717
	(0.213)	(0.030)	(0.722)	(0.436)
Observations	10,855	10,855	10,855	10,474
R^2	0.196	0.133	0.299	0.174
Stadium-FE	YES	YES	YES	YES
Season-FE	YES	YES	YES	YES
Referee-FE	YES	YES	YES	YES
r2_a	0.169	0.104	0.275	0.145
F	.	.	.	.
rss	25601	1791	238576	116946

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 8: Yellow Card Difference

VARIABLES	(1) Bundesliga	(2) La Liga	(3) Ligue 1	(4) Premier League	(5) Serie A
Fans Present	-0.435***	-0.650***	-0.543	-0.319***	-0.435***
	(0.143)	(0.166)	(0.454)	(0.111)	(0.102)
Average yellow cards earned by the home team	0.989***	0.987***	1.001***	0.997***	0.970***
	(0.021)	(0.015)	(0.038)	(0.021)	(0.018)
Average yellow cards given away by the home team	-0.997***	-0.971***	-0.968***	-0.979***	-0.979***
	(0.031)	(0.018)	(0.025)	(0.018)	(0.022)
Average yellow cards earned by the away team	-1.009***	-0.968***	-0.964***	-0.974***	-0.948***
	(0.103)	(0.088)	(0.107)	(0.096)	(0.101)
Average yellow cards given away by the away team	0.955***	0.976***	0.963***	0.964***	0.968***
	(0.140)	(0.091)	(0.091)	(0.087)	(0.075)
Constant	1.033***	0.929**	1.277***	0.523	0.033
	(0.354)	(0.448)	(0.454)	(0.317)	(0.312)
Observations	1,836	2,280	2,179	2,280	2,280
R^2	0.185	0.200	0.198	0.199	0.198
Stadium-FE	YES	YES	YES	YES	YES
Season-FE	YES	YES	YES	YES	YES
Referee-FE	YES	YES	YES	YES	YES
r2_a	0.153	0.174	0.169	0.174	0.161
F	.	.	.	.	.
rss	3855	6974	4466	4771	5479

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 6: Yellow Card Difference: Coefficient Plot of Dummy Variable for Fans

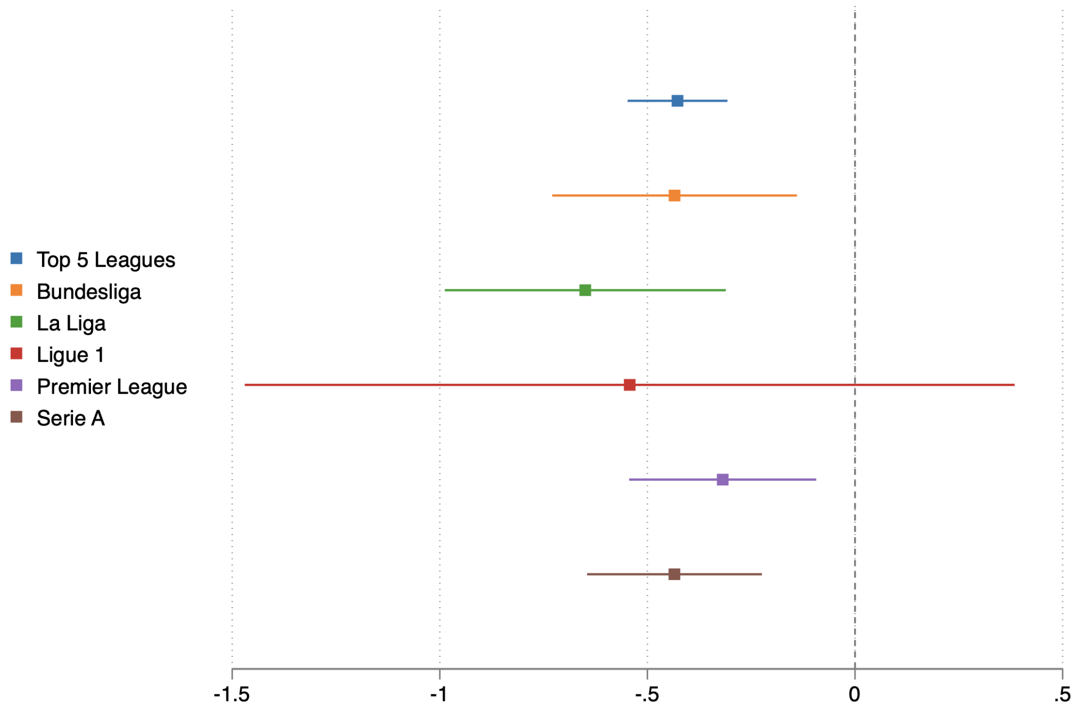

could vary and require their own regression. Table 8 shows the results of the regression in each league. Except for Ligue 1, the coefficients have significant p-values at the 1% level.

La Liga has the largest fan-induced difference while the Premier League has the smallest, with 0.65 and 0.319 more yellow cards awarded to the away team than the home team. These results show that fans cause the referee to give more yellow cards to the away team regardless of league. This is a new finding in the latest literature, where other papers found mixed results. With a larger sample size, I was able to show that fans have a consistent effect across the top 5 European Leagues.

6.3 Fans, the Referee, and Results

Given that there is evidence of a relationship between crowds and the referee and between crowds and final results, a case could be made that the referee may have a larger effect on match outcome than the fans.

To explore this, I ran a regression with the same controls as the initial regression in Section 6.1 and used a two-way fixed effect stadium and referee but included yellow card difference to go along with the Fans dummy variable.

Table 9: Goal Difference - Fans and Yellow Cards

VARIABLES	(1) Top 5 Leagues	(2) Bundesliga	(3) La Liga	(4) Ligue 1	(5) Premier League	(6) Serie A
Fans Present	0.243***	0.174	0.265***	0.465***	0.202	0.097
	(0.048)	(0.144)	(0.087)	(0.095)	(0.121)	(0.094)
Yellow Card Difference	-0.054***	-0.092***	-0.062***	-0.078***	-0.041	-0.014
	(0.010)	(0.025)	(0.019)	(0.022)	(0.025)	(0.024)
Constant	1.915***	3.248***	1.765**	-3.303***	1.931***	2.524***
	(0.396)	(1.112)	(0.853)	(1.067)	(0.697)	(0.811)
Observations	9,029	1,530	1,900	1,799	1,900	1,900
R^2	0.261	0.250	0.254	0.241	0.275	0.302
Stadium-FE	YES	YES	YES	YES	YES	YES
Referee-FE	YES	YES	YES	YES	YES	YES
r2_a	0.233	0.215	0.225	0.208	0.248	0.265
F	.	.	.	.	.	.
rss	22774	4541	4222	4463	5024	4403

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Goal Difference controls are included in the model but not displayed

Based on Table 9, both yellow card difference and the *Fans* coefficient are significant. Adding yellow card difference has increased fans' impact on the goal difference. The presence of fans results in an increase of 0.243 more home goals than away goals, just under a quarter of a goal more. In Ligue 1, fans are worth almost half a goal, 0.465 more home goals than away goals. This magnitude is substantial. It suggests that the mere presence of fans is almost halfway to a victory for the home team.

Looking at the *Yellow Card Difference* coefficients, in the Top 5 Leagues, an additional yellow card more for the home team is worth 0.054 goals for the away team. Intuitively, this result makes sense. An additional yellow card for the home team might cause them to play more cautiously and hold off on challenges more than they would have they weren't cautioned. This would slightly benefit the away team, as the coefficient shows.

These results suggest that fans are still a major reason for the observed home field advantage, even when accounting for referee specific variables. The mechanism goes directly through fans and not through fans to the referee and the match outcome.

6.4 Expected Goals

While there is evidence that the presence of a crowd results in an increase in goals for the home team, there are still questions about whether it changes the home team's performance. In a low-scoring game like soccer, one team can dominate proceedings but still lose the game by a one goal margin. One method that analytics in soccer has developed for measuring team performance is expected goals. Based on the cumulative probability of scoring chances, we can discern a more accurate representation of who deserved to win a match.

To test if there was an effect on the home team's performance, I ran a regression with the same controls as the Stadium Fixed Effects model in section 6.1 but with expected goal difference as the dependent variable.

In Table 10, we see that the effect of fans on expected goal difference is similar to the impact on goal difference. Across the Top 5 Leagues, the presence of the crowd results in 0.155 more home expected goals while holding all other variables constant. This implies that the crowd either facilitates the home team in attacking better or defending better. Running a subsequent regression on expected goals for and expected goals against, the improvement can be broken down into a 0.07 increase in expected goals for and a decrease of 0.086 expected goals against. Fans improve the underlying performance of the home team, with defense getting a slightly larger boost compared to attack.

In specific leagues, this effect varies in both magnitude and significance. In Germany, France, and Spain, Bundesliga, Ligue 1, La Liga, respectively, fans relate to an increase of between 0.188 and 0.249 more home expected goals than away expected goals. While in the Premier League and Serie A, fans are worth less than 0.1 more home expected goals; however, these results are insignificant.

6.4.1 Measuring Luck

An interesting application of expected goals is using them to measure "luck." The scoring luck a team experiences can be calculated by subtracting a team's expected goals from their actual goals for each match. The result is a team's overperformance (positive value) or underperformance (negative value). Adding this measure of luck to the stadium fixed effects regression will allow us to partially control for something previously immeasurable. As seen in Table 11, "luck" is the largest coefficient contributor to goal difference. Each additional goal over expectation for the home team, while holding other variables constant, results in an increase in goal difference of 0.909. Similarly, for the away luck coefficient, a goal over expectation for the away side results in a decrease in goal difference of 0.936. These coefficients are significant at the 1% level across all leagues. Essentially, the luckier a team is, the more likely they win, which is not a surprise.

Table 10: Expected Goals

VARIABLES	(1) Top 5 Leagues	(2) Bundesliga	(3) La Liga	(4) Ligue 1	(5) Premier League	(6) Serie A
Fans Present	0.155***	0.189**	0.188***	0.249***	0.084	0.076
	(0.032)	(0.080)	(0.049)	(0.087)	(0.087)	(0.057)
Home Soccer Power Index	0.020***	0.016*	0.018**	0.018***	0.020***	0.021***
	(0.003)	(0.008)	(0.009)	(0.007)	(0.006)	(0.007)
Away Soccer Power Index	-0.042***	-0.045***	-0.040***	-0.038***	-0.044***	-0.041***
	(0.001)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)
Distance in 100KM	0.008	0.006	0.016**	-0.009	0.003	0.017*
	(0.005)	(0.024)	(0.008)	(0.012)	(0.024)	(0.008)
Home Rest	-0.001	-0.003	-0.003	-0.000	0.002	-0.000
	(0.001)	(0.003)	(0.003)	(0.003)	(0.002)	(0.001)
Away Rest	0.001	0.004	0.001	0.000	-0.001	0.001
	(0.001)	(0.004)	(0.003)	(0.004)	(0.003)	(0.002)
Home Team in the Champions League	-0.024	-0.101	-0.211**	-0.090	0.117	0.101
	(0.057)	(0.074)	(0.079)	(0.155)	(0.117)	(0.129)
Away Team in the Champions League	-0.153***	-0.099	-0.142*	-0.498***	-0.058	-0.081
	(0.039)	(0.089)	(0.078)	(0.116)	(0.072)	(0.081)
Home Team in the Europa League	0.077*	0.006	-0.003	0.227**	0.225***	0.030
	(0.046)	(0.118)	(0.083)	(0.087)	(0.069)	(0.111)
Away Team in the Europa League	-0.019	-0.037	0.011	-0.065	0.098	-0.112
	(0.035)	(0.095)	(0.074)	(0.073)	(0.074)	(0.069)
VAR Implemented	-0.008	-0.095	0.022	0.037	-0.008	-0.041
	(0.034)	(0.090)	(0.063)	(0.063)	(0.105)	(0.072)
Constant	1.612***	2.220***	1.753**	1.337***	1.850***	1.467***
	(0.220)	(0.725)	(0.648)	(0.437)	(0.402)	(0.420)
Observations	9,029	1,530	1,900	1,799	1,900	1,900
R^2	0.349	0.329	0.319	0.317	0.381	0.396
Stadium-FE	YES	YES	YES	YES	YES	YES
Season-FE	NO	NO	NO	NO	NO	NO
r2_a	0.337	0.313	0.304	0.302	0.367	0.384
F	212.7	86.51	76.45	90.29	83.90	105.6
rss	8460	1631	1525	1719	1823	1720

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The fans coefficient with the luck model shows a slight decrease in the crowds' impact. While statistically significant in three of the five leagues, the coefficient has dropped to only 0.16 more home goals than away goals overall. Luck appears to have a more considerable impact on match outcome than fans do, a finding that intuitively is not surprising. However, there are some limitations to these findings and this measure of "luck." Most notably, home and away team luck on an individual scale can say more about a team's ability than about luck. Expected goal metrics are not player specific; a very talented goalscorer, like Lionel Messi, will over perform their expected goals throughout a season. Teams that consistently overperform their expected goals will generally be better and win more.

Table 11: Goal Difference with Luck

VARIABLES	(1) Top 5 Leagues	(2) Bundesliga	(3) La Liga	(4) Ligue 1	(5) Premier League	(6) Serie A
Fans Present	0.159***	0.183**	0.195***	0.266***	0.091	0.071
	(0.032)	(0.081)	(0.048)	(0.087)	(0.085)	(0.056)
Home Team Luck	0.909***	0.930***	0.916***	0.891***	0.933***	0.879***
	(0.011)	(0.030)	(0.015)	(0.023)	(0.024)	(0.025)
Away Team Luck	-0.936***	-0.975***	-0.923***	-0.947***	-0.965***	-0.870***
	(0.012)	(0.026)	(0.031)	(0.027)	(0.027)	(0.021)
Constant	1.625***	2.282***	1.781***	1.271**	1.868***	1.483***
	(0.220)	(0.728)	(0.641)	(0.464)	(0.401)	(0.432)
Observations	9,029	1,530	1,900	1,799	1,900	1,900
R^2	0.729	0.732	0.734	0.712	0.738	0.737
Stadium-FE	YES	YES	YES	YES	YES	YES
r2_a	0.725	0.725	0.728	0.706	0.732	0.731
F	1580	1338	1261	279.0	610.6	533.6
rss	8347	1620	1503	1692	1812	1661

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Goal Difference controls are included in the model but not displayed

Additionally, my measure of luck is closely related to goal difference, the two luck variables are the two variables most correlated with goal difference. This could create some bias, especially because both are post-match performance measures. A team with a positive goal difference will also likely have positive luck, with the two being directly linked as post-match performance measures.

6.5 Robustness Check

To test the robustness of my findings (Table 12), I changed the variable of interest to Percent Full, a measure of stadium density and explored home field advantage using structural break graphs. Even with the new variable and the increased number of values (rather than just 1 or 0), little has changed. For each additional percent increase in *Percent Full*, holding all else constant, the home team is expected to score 0.002 more goals. The value with a full stadium (*Percent Full* = 100) is very similar to the value with the dummy variable, 0.2 more goals for the home team. Similarly, with yellow card difference, for each additional percent, the referee awards 0.003 more yellow cards to the away team compared to the home team. This value is smaller than the number found with the dummy variable but still statistically significant.

Table 12: Robustness Check - Percent Fans as Predictor of Interest

	(1)	(2)
VARIABLES	Goal Difference	Yellow Card Difference
Percent Full	0.002***	-0.003***
	(0.001)	(0.001)
Constant	1.913***	0.018
	(0.379)	(0.168)
Observations	9,029	10,475
R^2	0.244	0.102
Stadium-FE	YES	YES
Season-FE	NO	NO
r2_a	0.231	0.0887
F	103.5	65.08
rss	23324	27648

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Controls are included but not shown

6.5.1 Structural Break Graphs

The COVID interruption in March 2020 created a structural break in the data. The immediate shock of playing matches behind closed doors was emotionally significant. Splitting the data into two groups, pre-COVID and post-COVID, will help explain some of what happened during this period.

As evidenced in Figure 7, which depicts the home team winning percentage (home wins - home losses and excluding draws divided by the number of matches played each date) on each match day[7], the slope in the two groups is quite different, there is a slight increasing relationship between home team win percentage and date before COVID but after the resumption, the slope increases quickly back to pre-pandemic levels. There is also a notable drop in intercept, with the intercept of winning percentage dropping down to 50% before increasing back. This lends some credence to the hypothesis that players may have gotten used to the lack of fans, and their performance levels increased back to normal after the initial shock wore off.

One critique on structural breaks and regression discontinuities is assigning a large weight to where we place the break, placing the break at the COVID resumption is logical but the findings could be purely coincidental. In Figure 7, it looks like the COVID imposed break had a large impact on the home winning percentage. However, when I arbitrarily assign a break before the 2019 season, we see a different change but one that looks considerable in magnitude. Like above, the break in Figure 8 shows that there is a small jump between the two groups but the slope changes from an increasing relationship

[7]If there are three matches played today, the home wins one, loses one, and draws one, the home team winning percentage is 50%. The draw is excluded and the home team has won one out of two matches.

Figure 7: Structural Break of Home Winning Percent on a Match Day - Line denotes COVID

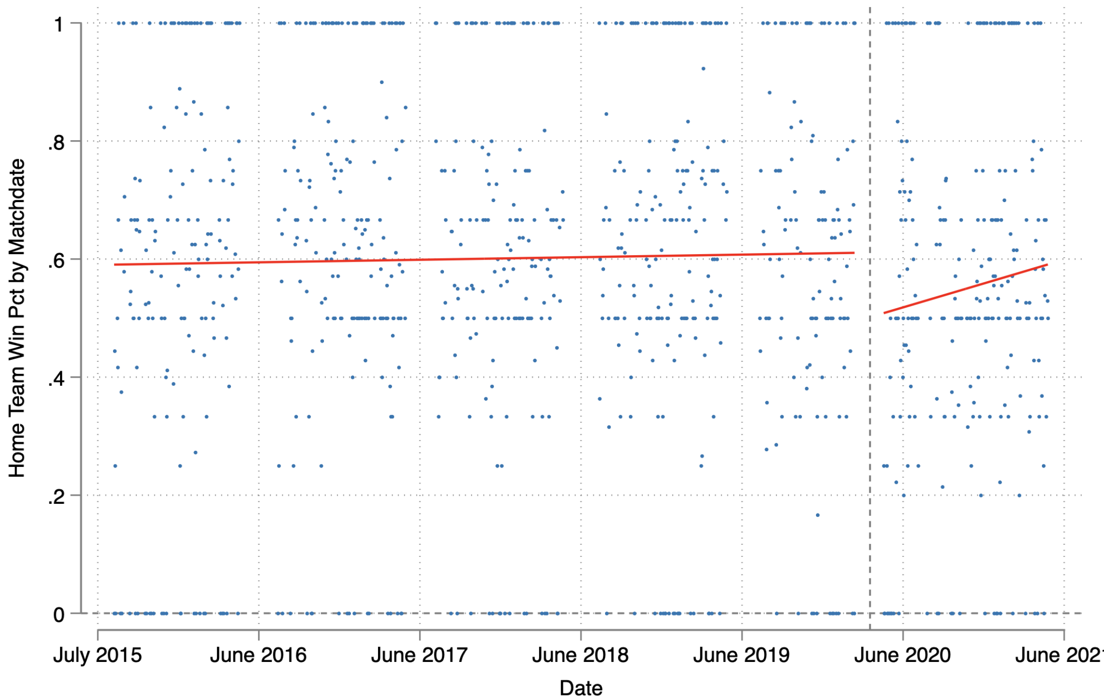

to a decreasing one. Looking at this Figure 8 alone the relationship between home winning percentage and date is dubious. However, an Chow F Test on the significance of the interaction variable and the COVID dummy variable (F Test = 0.0494) means that there is slight evidence that the structural break is significant.

7 Discussion

My paper utilizes a fixed effects regression to study crowds' effect on two soccer-related outcomes, the match, and the referee. I capitalize on the quasi-natural experiment that the COVID-19 pandemic created to study the impacts of fans' absence, using data from both COVID-affected seasons, the latter half of the 2019 season, and the entire 2020 season. This extra season of data means that this paper is one of the first to analyze the effects of home field advantage in the French Ligue 1, which canceled the remainder of the 2019 season.

Figure 8: Structural Break of Home Winning Percent on a Match Day using a Placebo

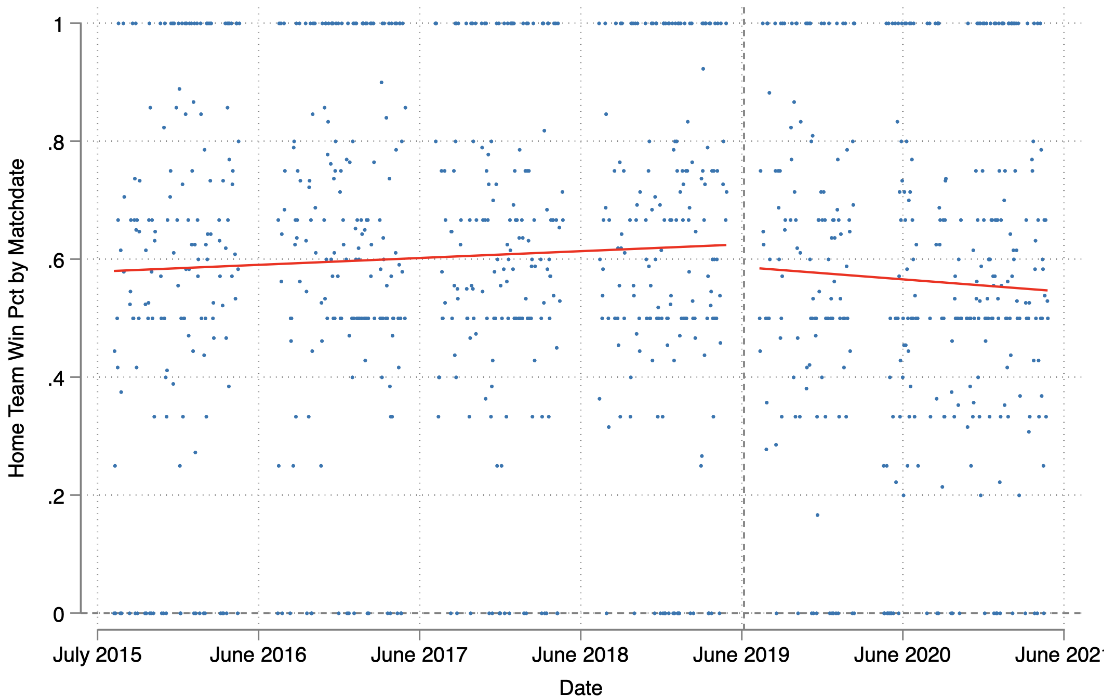

Evidence from the top 5 European Leagues consistently supports the hypothesis that the lack of fans reduced home field advantage, and their presence results in a positive gain for the home team. This is evident in both the regressions with goal difference. Overall, fans were worth 0.2 more goals, and the regressions with the selected referee variables, where fans all increased the difference between the home and away team in favor of the home team. However, these effects did vary within league. The presence of fans in the Bundesliga is only worth 0.09 more home goals, a surprise considering they were a league that was cited as experiencing a sharp initial decline in home field advantage in papers that study the second half of the 2019 season (Benz & Lopez, 2021). Similarly, the effect in the Serie A was muted, with a mere 0.03 more goals. However, in the other three leagues, the evidence that home field advantage increases with fans was evident. Ligue 1 has the largest home field advantage, fans are worth just under half of a goal more, a very meaningful value in a low scoring game such as soccer. This implies that in a match between two equal teams, the presence of fans is almost halfway to a victory for the home team (a goal difference of 1 would be a home team win).

A similar trend is apparent regarding yellow card difference within each league, although the coefficients are more consistent between the leagues. All the coefficients agree that there is a negative relationship between the crowd and the referee. The crowd influences the referee into giving more yellow cards to the away team, ranging in magnitude from 0.32 in the Premier League to 0.65 in La Liga. These results show strong evidence of home field advantage in the spread of referee decisions.

It is worth noting that these yellow card differences do not necessarily significantly impact match outcome, a conclusion that Benz and Lopez (2021) also arrive at. In the Bundesliga and Serie A, there was weak evidence of home field advantage as it pertains to the match outcome, but both leagues had strong evidence of home field advantage in yellow card difference. This suggests yellow cards and goal difference are not linked. This could be explained by Carmichael and Thomas's (2005) findings that the team's tactics changed depending on whether they were home or away. Home teams could have more possession, causing the away team to defend more and, therefore, more opportunities to commit fouls. Regardless of whether there is a link between yellow cards and goal difference, it is apparent that fans influence the referee.

Using the Bundesliga as another example, the previous findings of home advantage in this league differ from the findings above. One possible explanation is that the initial shock of playing in an empty stadium wore off. The lack of fans was no longer salient, and home field advantage returned, albeit without fans. This is clear in Figure 3, which shows that the Bundesliga and Serie A were the only two leagues that saw an increase in home field advantage during the 2020 season, after a steep fall during the 2019 season. A possible explanation is that teams in these leagues were initially affected by the lack of fans but slowly adjusted. In contrast, teams in other leagues gradually felt the effect, with the reduced generalized drive affecting them more in the 2020 season.

One limitation is that this paper does not only study the effects of the lack of fans but also the new world players had to adapt to with the COVID-19 pandemic. These results say just as much about the new pandemic conditions as they do about the cause and effect from the lack of fans. Changes in training methods or players' lives could also contribute to the decrease in home field advantage observed across the leagues. Leagues around the world adapted to make playing during the pandemic easier. One of these major changes was increasing the number of substitutes from 3 to 5 in all of the leagues. This would favor bigger and richer teams, as they would have more depth. Changes like this could impact the observed decrease in home field advantage.

There is a subsequent question that I still want to answer in future research, though it comes with its empirical challenges. One thing that might be missing from this analysis and the assumption that empty stadiums have a negative effect on players' performance are cases where pre-pandemic teams already played in front of small crowds. Did these teams adjust more quickly to the post-pandemic empty stadiums because it was closer to the norm? One major confounder of this question is the link between team quality and attendance. Teams that frequently play in front of small crowds are usually weaker than their opposition, so they might frequently lose as the home team.

8 Conclusion

This paper shows that the lack of fans during the COVID-19 pandemic reduced home field advantage in the top 5 European Leagues. While there is not enough statistical evidence to reject the null hypothesis that fans do not affect either goal difference or the referee for all the league-specific variables, there is evidence to reject the null hypothesis on most league-specific variables, both in goal and yellow card difference. I can also reject the null hypothesis on goal difference and the referee decision variables when using the entire dataset. These conclusions support the argument that the crowd is a substantial driver in home field advantage.

As mentioned above, these results, while significant, should be conditioned by the fact that teams had multiple stressors to adapt to, rather than just the lack of fans. During the pandemic, the top 5 European Leagues all experienced decreases in home advantage in terms of goal difference, home point to away point ratio, and the number of referee decisions.

9 Bibliography

References

Anders, A., & Rotthoff, K. W. (2014). Is home-field advantage driven by the fans? evidence from across the ocean. *Applied Economics Letters*, *21*(16), 1165–1168. https://doi.org/10.1080/13504851.2014.914139

Arteta, M. (2021, August 13). 'fans are the club's soul, they offer us energy'. Retrieved March 6, 2022, from https://www.arsenal.com/news/fans-are-clubs-soul-they-offer-us-energy

Benz, L. S., & Lopez, M. J. (2021). Estimating the change in soccer's home advantage during the covid-19 pandemic using bivariate poisson regression. *arXiv:2012.14949 [stat]*. Retrieved September 12, 2021, from http://arxiv.org/abs/2012.14949

Bond, C., & Titus, L. (1983). Social facilitation: A meta-analysis of 241 studies. *Psychological bulletin*, *94*, 265–92. https://doi.org/10.1037/0033-2909.94.2.265

Boyko, R. H., Boyko, A. R., & Boyko, M. G. (2007). Referee bias contributes to home advantage in english premiership football. *Journal of Sports Sciences*, *25*(11), 1185–1194. https://doi.org/10.1080/02640410601038576

Buraimo, B., Forrest, D., & Simmons, R. (2010). The 12th man?: Refereeing bias in english and german soccer. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *173*(2), 431–449. https://doi.org/10.1111/j.1467-985X.2009.00604.x

Carmichael, F., & Thomas, D. (2005). Home field effect and team performance: Evidence from english premiership football. *Journal of Sports Economics*, *6*(3), 264–281. https://doi.org/10.1177/1527002504266154

Clarke, S. R., & Norman, J. M. (1995). Home ground advantage of individual clubs in english soccer. *Royal Statistical Society Journal series D: The Statistician*, *44*, 509–521. Retrieved October 2, 2020, from https://researchbank.swinburne.edu.au/items/4b4ca583-30a6-4dc2-adc6-8848b6f095cd/1/

Davis, J. L., & Krieger, K. (2021, August 30). *How much of 'home field advantage' comes from the fans? a natural experiment from the COVID-19 pandemic* (SSRN Scholarly Paper ID 3914238). Social Science Research Network. Rochester, NY. https://doi.org/10.2139/ssrn.3914238

Edwards, J. D., & Archambault, D. (1988). The home-field advantage. *Sports, games, and play* (2nd Edition, p. 38). Psychology Press. Retrieved October 4, 2021, from https://www.taylorfrancis.com/chapters/edit/10.4324/9780203728376-17/home-field-advantage-john-edwards-denise-archambault

Fischer, K., & Haucap, J. (2020). *Does crowd support drive the home advantage in professional soccer? evidence from german ghost games during the covid-19 pandemic* (SSRN Scholarly Paper ID 3697062). Social Science Research Network. Rochester, NY. Retrieved September 28, 2021, from https://papers.ssrn.com/abstract=3697062

Hegarty, T. (2021). Information and price efficiency in the absence of home crowd advantage. *Applied Economics Letters*, *0*(0), 1–6. https://doi.org/10.1080/13504851.2021.1883525

Hill, Y., & Van Yperen, N. W. (2021). Losing the home field advantage when playing behind closed doors during COVID-19: Change or chance? *Frontiers in Psychology*, *12*, 1321. https://doi.org/10.3389/fpsyg.2021.658452

Johnston, R. (2008). On referee bias, crowd size, and home advantage in the english soccer premiership. *Journal of Sports Sciences*, *26*(6), 563–568. https://doi.org/10.1080/02640410701736780

McCullagh, P. D., & Landers, D. M. (1976). Size of audience and social facilitation [Publisher: SAGE Publications Inc]. *Perceptual and Motor Skills*, *42*(3), 1067–1070. https://doi.org/10.2466/pms.1976.42.3c.1067

Meissner, S. (1994). Components of social facilitation, home field advantage, and home field disadvantage and their effects on sport performance. *Graduate Student Theses, Dissertations, & Professional Papers*. https://scholarworks.umt.edu/etd/1876

Nevill, A., Balmer, N., & Mark Williams, A. (2002). The influence of crowd noise and experience upon refereeing decisions in football. *Psychology of Sport and Exercise*, *3*(4), 261–272. https://doi.org/10.1016/S1469-0292(01)00033-4

Olley, J. (2020, June 30). *Premier league games with no fans: How players, coaches cope with the silence* [ESPN.com]. Retrieved December 2, 2021, from https://www.espn.com/soccer/english-premier-league/story/4122080/premier-league-games-with-no-fans-how-playerscoaches-cope-with-the-silence

Pollard, R. (1986). Home advantage in soccer: A retrospective analysis. *Journal of Sports Sciences*, *4*(3), 237–248. https://doi.org/10.1080/02640418608732122

Pollard, R., & Pollard, G. (2005). Home advantage in soccer: A review of its existence and causes. *International Journal of Soccer and Science*, *3*(1), 25.

Ponzo, M., & Scoppa, V. (2018). Does the home advantage depend on crowd support? evidence from same-stadium derbies. *Journal of Sports Economics*, *19*(4), 562–582. https://doi.org/10.1177/1527002516665794

Reade, J. J., & Koyama, M. (2009). Playing like the home team: An economic investigation into home advantage in football. *International Journal of Sport Finance*, *4*, 16–41.

Reade, J. J., Schreyer, D., & Singleton, C. (2020). Echoes: What happens when football is played behind closed doors? *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3630130

Triplett, N. (1898). The dynamogenic factors in pacemaking and competition. *The American Journal of Psychology*, *9*(4), 507–533. https://doi.org/10.2307/1412188

Zajonc, R. B. (1965). Social facilitation. *Science*, *149*, 269–274. https://doi.org/10.1126/science.149.3681.269