

Games of Two Halves: Non-Experimental Evidence on Cooperation, Defection and the Prisoner's Dilemma

STEPHEN DOBSON

*Hull University Business School, University of Hull **

JOHN GODDARD

Bangor Business School, Bangor University

We develop a stylized two-period game-theoretic model of the strategic choices made by soccer teams when selecting between defensive and attacking team formations, and between non-violent and violent styles of play. Cooperative behaviour during the early stages of matches is typically superseded by non-cooperation during the latter stages. The propensity for violent play to take place in the latter stages of soccer matches is interpreted as novel non-experimental evidence that players typically resort to mutually detrimental non-cooperative forms of behaviour when the payoffs assume a prisoner's dilemma structure.

Keywords: Game theory, Prisoner's dilemma, Hazard functions

JEL Classifications: C41, C70, L83

1 Introduction

In soccer (football), it is well known that the incidence of player dismissals for acts of serious foul play, signalled by the brandishing of a 'red card' by the match referee, is higher during the latter stages of matches than during the early stages. For example, the Merseyside 'derby' Everton *versus* Liverpool, one of the most fiercely contested fixtures in the English Premier League calendar, has been called "the friendly derby" in acknowledgement of the camaraderie that exists between the two sets of supporters. The historical record, however, reveals that this fixture has produced more red cards than any other since the inception of the Premier League in the 1992/93 season: 21 in 51 league meetings to the end of 2017.¹ Only three of these 21

* Corresponding author: Professor Stephen Dobson, Hull University Business School, University of Hull, Cottingham Road, Hull, HU6 7RX, UK, s.m.dobson@hull.ac.uk

© 2018 Stephen Dobson and John Goddard. Licensed under the Creative Commons Attribution - Noncommercial 3.0 Licence (<http://creativecommons.org/licenses/by-nc/3.0/>). Available at <http://rofea.org>.

¹ See 'Liverpool against Everton is the friendly derby but there's been more red cards in this fixture than any other in the Premier League', Daily Mail 19/12/2016, <http://www.dailymail.co.uk/sport/football/article-4047792.html>

dismissals took place in the first halves of these matches; the remaining 18 occurred in the second halves; and no fewer than 12 of these during the final 15 minutes of regular play.

Football pundits typically attribute such a pattern in the timing of red cards to player fatigue, and as a consequence an increased incidence of errors such as mistimed tackles, in the latter stages of matches. Although we accept that the player fatigue hypothesis may form part of the explanation for the observed pattern, in this paper we develop an alternative explanation based on axioms of rational, optimizing strategic behaviour. Adopting a game-theoretic approach, we argue that it is common to observe both cooperative and non-cooperative forms of behaviour at different stages within the 90-minute duration of a single match. In the present context, cooperation involves playing within the laws of the sport, in a non-violent style. Non-cooperation involves the adoption of a violent style of play, including tackling in a manner that is proscribed by the laws of the sport² and other unlawful forms of violent conduct or sabotage that may be intended to gain an unfair advantage by intimidating or causing injury to opposing players, or otherwise disrupting the regular flow of play

The idea that cooperation can be mutually advantageous, but is liable to break down when the structure of rewards and costs offers incentives for individual players to behave in a non-cooperative fashion, is a core result in game theory. The prisoner's dilemma has been used extensively as a pedagogic device, to illustrate the notion that actions that appear to be desirable from an individual player's perspective may produce outcomes that are seen to be undesirable when the players' interdependence is considered. The resulting Pareto-suboptimal Nash equilibrium has been used to describe the outcomes of competitive games in a wide range of spheres, including the economic, social, political, military and biological.

To explain the pattern in the timing of red cards, we develop a game-theoretic interpretation of the strategic choices made by soccer teams when selecting between defensive and attacking team formations, and between non-violent and violent styles of play. We demonstrate that the structure of the rewards and costs to soccer teams is such as to encourage cooperation (a non-violent style of play) during the early stages of matches, often followed by defection to non-cooperative (violent) behaviour during the latter stages. The rewards and costs of violence in the latter stages of matches that remain level in goals scored exhibit a prisoner's dilemma structure.

Our analysis is based on the observation that the structure of payoffs for a soccer match is non-zero sum. This is true in the trivial sense that three league points are awarded for a win, but only one point to each team for a draw. We focus, however, on another non-zero sum aspect of the payoff structure, which arises from the punishment that is imposed when a player is dismissed for violent play, with no replacement being permitted. A team that incurs

² For example, tackling an opponent from behind, tackling using both feet, and tackling in a reckless manner that is likely to cause injury, are all proscribed modes of tackling.

a red card is penalized in two ways: first, through completing the present match with a reduced complement of players; and second, through the suspension of the dismissed player from a number of future matches. The team whose player has incurred a red card is permitted to field a full complement of 11 players in the matches from which the player is suspended, but the 11 are selected from a depleted and weakened squad.³ For a player who has not been dismissed previously within the current soccer season, the duration of the suspension is between one and three future matches, depending on the nature of the red-card offence. In the case of a player who has been dismissed previously within the current season, the duration of the suspension is increased by one match for each previous red card.

Accordingly, the total cost to the team that incurs a red card comprises: (i) a reduction in the expected number of league points gained from the present match; and (ii) a reduction in the expected league points gained from future matches from which the dismissed player is suspended. The impact of (i) on the payoffs for the present match is zero-sum (the opposing team gains as the penalized team loses); but the impact of (ii) is non-zero sum because the cost to the penalized team is not counterbalanced by any equivalent gain to the opposing team in the present match.

The magnitude of (i) is dependent on the timing of the red card within the present match: a red card incurred in the early stages imposes a larger cost, because the team plays at a numerical disadvantage for a longer period. By contrast, the magnitude of (ii) is independent of the timing of the red card within the match. The decline in the magnitude of (i) over the match duration creates a prisoner's dilemma payoff structure during the latter stages of matches that remain level in goals scored, and a tendency for cooperation during the early stages of matches to be superseded by non-cooperative behaviour during the latter stages.

Before developing the game-theoretic model, we present a more detailed empirical investigation that reveals several interesting patterns in the timings of goals and red cards, conditional on the current state of the match (difference in goals scored). We observe, for example, that the scoring rates of teams that are currently trailing in goals scored, after controlling for team quality, are significantly higher than those of teams that are leading or level in scores. Furthermore, teams that are currently trailing experience an increased probability of incurring a red card, while teams that are currently leading appear to play more defensively and less violently. These patterns cannot be explained by the player fatigue hypothesis, but they are consistent with the strategic approach that is captured by the theoretical model.

The main caveat associated with the empirical investigation is that only the outcomes (red cards and goals) that determine the teams' payoffs are observed, while the teams' strategic

³ The dismissed player may incur also costs in the form of reputational damage, or the risk that a first-team place lost during a suspension may not be recovered when the suspension is complete.

choices are unobserved. Accordingly, we do not present empirical evidence in the (ideal) form of direct observation of strategies. This limitation is offset, however, by a major strength of our non-experimental set-up. The aforementioned shift in the configurations of rewards and costs over the duration of the match, and the emergence of a sub-game with a prisoner's dilemma structure in the latter stages, permits the observation of outcomes from cooperative and non-cooperative forms of behaviour on the part of the same subjects (the two opposing teams) under constant environmental conditions (stadium, pitch, spectators, weather and match officials). Accordingly, and in contrast to all previous non-experimental studies of which we are aware, we are able to examine the effects on subject behaviour of changes in the payoff structure in isolation from the effects of variations in subject characteristics and environmental conditions.

The paper also investigates the kind of strategies that are consistent with the patterns in outcomes reported in the empirical analysis. Consequently, in the theoretical model, a soccer team that adopts an attacking formation increases its own probability of scoring a goal, but it also increases its opponent's scoring probability and its own probability of incurring a red card for committing a last-ditch foul in order to prevent an opposing player from scoring. A team that plays in a violent style engages in sabotage by operating outside the laws of the game and disrupting the regular flow of play, at a heightened risk that one of its own players will incur a red card. The disruptive aspect of violent play increases the probabilities of goals being scored by either team: errors are more likely to be committed, leading to goals being scored, when play is reckless and tempers become frayed.

Since the rewards from adopting an attacking formation or playing in a violent style take the form of an increase in the team's own scoring probability, the incentive to attack or play violently is greatest for a team that is currently trailing in goals scored. The costs take the form of an increase in the opposing team's scoring probability, and an increase in the team's own probability of incurring a red card. As we have noted above, a red card places the team at a numerical disadvantage for the remainder of the present match, and weakens the team in future matches from which the dismissed player is suspended (because the team will be selected from a depleted squad).

The economics literature presents extensive experimental evidence on the adoption of cooperative and non-cooperative behaviour; but non-experimental evidence is, understandably, sparse. Experimental evidence covers finitely-repeated games (Andreoni and Miller, 1993; Bereby-Myer and Roth, 2006; Mengel, 2014; Kagel and McGee, 2016; Embrey et al, 2017; Mao et al, 2017), infinitely-repeated games (Dal Bó, 2005; Duffy and Ochs, 2009; Dal Bó and Fréchette, 2011, 2015; Fréchette and Yuksel, 2017) and repeated games played in continuous time (Friedman and Oprea, 2010; Bigoni et al, 2015). Typically, experiments are based on a series of predefined games. Estimation of the parameters of interest is based on the observation of subjects' behaviour (for a recent example, see Lugovskyy et al, 2017).

Drawbacks of experimental settings include the selectivity of participants (usually college students), the limited realism of the laboratory environment, the high level of contextual information provided to participants, and the unusual degree of scrutiny (Levitt and List, 2007). In an attempt to overcome these limitations, prisoner's dilemma games that arise in non-experimental settings have received attention in the literature. Levitt (2004), List (2006), Belot et al. (2010), Oberholzer-Gee et al. (2010), Van den Assem et al. (2011) and Burton-Chellew and West (2012) examine data from TV game shows. An advantage of the game show setting, in contrast to many experiments, is that the decisions of participants are observed publicly. This adds a sense of realism, since decisions are rarely made anonymously in real life. However, public observation may tend to amplify cooperation, if a selfish participant chooses to cooperate only in order to avoid embarrassment. Another drawback is that game show participants are not representative of the population, since they are not randomly selected.⁴

The rest of this study proceeds as follows. Section 2 reports estimation results for an empirical model for the in-play arrival rates of goals and red cards (player dismissals) in more than 18,000 professional matches played in the English leagues during the period 2001-2010. Section 3 develops a stylized two-period game theoretic model. The two periods correspond to the first and second halves of a regular 90-minute soccer match. Conditions are derived for the emergence of a Pareto-optimal equilibrium in the first half, and a Pareto-suboptimal Nash equilibrium with a prisoner's dilemma structure in the second half. These conditions are satisfied by reasonable 'ballpark' parameter estimates. Section 4 summarizes and concludes.

2 Empirical Investigation

The empirical investigation seeks to identify variations in the arrival rates of red cards and goals over the duration of the match, conditional on the state of the match (defined, as before, by the difference in goals already scored and the numbers of players already having been awarded a red card and dismissed). The data sample for the empirical analysis that is reported in this study comprises all 18,324 matches played in the four tiers (divisions) of the English Premier League and Football League during the nine soccer seasons from 2001/02 to 2009/10 (inclusive). As noted previously, we do not observe the teams' strategic choices (defend or attack, non-violent or violent) directly, but we observe realizations of goals scored and red

⁴ Participation in professional soccer is also highly selective, and public scrutiny of the professional sport is intense. As a source of non-experimental evidence, therefore, professional soccer is subject to some limitations that are similar to those of TV game shows. TV game show participants, however, represent a tiny proportion of the population; by contrast, several million people in England have direct personal experience of participation in soccer at a non-professional level, either currently or in the past. According to Sport England's Active People Survey, estimated participation in soccer at least once per month in England was 3.1 million in 2007, and 2.6 million in 2016.

cards; and it is reasonable to expect that these observable metrics correlate with the teams' chosen strategies.⁵

Following a modelling approach similar to that of Dixon and Robinson (1998), we assume that the arrival rates of red cards and goals can be represented as Poisson processes, such that the probability that a new arrival occurs is independent of the time that has elapsed since the previous arrival. It is also assumed, for simplicity, that arrivals of red cards and goals are independent.⁶ Let μ_k denote the arrival rate for event k , where $k=1$ denotes a home-team red card, $k=2$ denotes an away-team red card, $k=3$ denotes a goal scored by the home team, and $k=4$ denotes a goal scored by the away team. Let m_k denote the number of minutes that elapse before the next occurrence of event k . If red cards and goals are Poisson processes, m_k follows an exponential distribution, with distribution function $F_k(t) = \text{prob}(m_k \leq t) = 1 - \exp(-\mu_k t)$ and density function $f_k(t) = F_k'(t) = \mu_k \exp(-\mu_k t)$.

We estimate a competing-risks model, comprising hazard functions for the four events ($k=1, \dots, 4$), with μ_k assumed to be linear in a set of covariates that are either constant or time-varying over the duration of the present match. Spells of continuous play that end in the occurrence of event k are treated as right-censored in the likelihood functions for the events other than k . The full list of covariate definitions for the conditional hazard functions for home- and away-team red cards and goals is as follows:

RELQUAL = HPROB + 0.5 × DPROB, where HPROB and DPROB are implicit prior probabilities (before the start of the match) of a home win and draw, respectively, computed from data on prices for fixed-odds betting posted by a panel of high-street and internet bookmakers.⁷

DUR = duration (number of minutes elapsed prior to the start of the current minute) within the present match, measured from 0 to 89.

⁵ Garicano and Palacios-Huerta (2005) report empirical evidence of the effect of a change of incentives, arising from an increase in the number of league points awarded for a win from two to three, on a range of metrics for Spanish soccer, including goals, yellow and red cards, fouls and shots.

⁶ This is a simplification because a player who commits a foul in his own team's penalty area, and by so doing denies the opposing team a goal-scoring opportunity, is liable to incur a red card, while the opposing team has an opportunity to score from the resulting penalty kick. The proportions of all red cards and goals that arrive simultaneously in this manner are small, however. Ignoring this element of simultaneity greatly simplifies the specification and estimation of the empirical model.

⁷ Prices for fixed-odds betting were obtained from www.football-data.co.uk for the bookmakers Interwetten, Ladbrokes and William Hill. Let θ_j denote the decimal odds for bets on a home win ($j=1$), draw ($j=2$) and away win ($j=3$). The implied probabilities for the three outcomes are $\lambda_j = \theta_j^{-1} / \sum_k \theta_k^{-1}$. HPROB and DPROB are the arithmetic means of λ_1 and λ_2 , respectively, calculated over the three sets of bookmaker odds (or those odds that were available in the case of a few matches for which the bookmaker data were incomplete).

UNCERT = RELQUAL \times (1-RELQUAL).

M45, M90 = 0-1 dummy variables that allow for step changes in the arrival rates recorded for the 45th and 90th minutes of the match, due to stoppage time.

KICKOFF = 1 for the 1st and 46th minutes, and for any minute when a goal has been scored by either team in the preceding minute, and 0 otherwise.

DIFFs = 0-1 dummy variables indicating the goal difference between the home and away teams at the start of the current minute: DIFF-3 = 1 if the away team is leading by three goals or more at the start of the current minute, and 0 otherwise. DIFF-2 = 1 and DIFF-1 = 1 if the away team is leading by two goals or one goal, respectively. DIFF+1, DIFF+2 and DIFF+3 indicate the home team leading by one goal, by two goals, or by three goals or more, respectively.

HOFF, AOFF = 0-1 dummy variables indicating whether either team is currently at a numerical disadvantage owing to one or more players having incurred a red card previously during the present match. HOFF = 1 if the home team is at a disadvantage in player numbers, and 0 otherwise. AOFF is defined similarly for the away team.⁸

By the regulations, a player may be dismissed for committing either a serious offence, punishable by the award of a red card, or two less serious offences, punishable by the award of a yellow card (caution), within the present match. The latter creates a complication for the interpretation of the data on red cards conditional on duration: an incidence of yellow cards that was independent of duration would create a spurious pattern of duration dependence in the data on red cards. In the data set, 1,433 of the 3,396 red cards (42.2 per cent) were awarded for a second yellow-card offence. Table 1 reports two sets of red card hazard function estimates: for all red cards in columns (1) and (2); and for red cards excluding those awarded for two yellow-card offences in columns (3) and (4). The hazards for goals scored by the home and away teams are reported in columns (5) and (6).

2.1 Red Card Hazard Functions

In the hazard functions for all home- and away-team red cards, the coefficients on M45 and M90 reflect the practice of recording any event that occurs during stoppage time at the end of each 45-minute period as having occurred in either the 45th or the 90th minute. The hazards

⁸ The endogenous nature of several right-hand-side variables in the hazard functions is not considered in the estimations. The variables KICKOFF, DIFFs, HOFF and AOFF are all dependent on previous arrivals of red cards and goals. However, we are unaware of any empirical specification that would accommodate endogeneity of this form.

Table 1: Estimated Hazard Functions for Red Cards and Goals

	All red cards		Red cards excluding two yellow cards		All goals	
	Home team (1)	Away team (2)	Home team (3)	Away team (4)	Home team (5)	Away team (6)
KICKOFF	-	-	-	-	-.6204 ^{***} (-15.8)	-.3294 ^{***} (-8.33)
M45	1.9163 ^{***} (15.3)	1.4322 ^{***} (11.6)	2.0207 ^{***} (13.9)	1.3956 ^{***} (9.03)	1.2625 ^{***} (38.7)	1.2144 ^{***} (31.5)
M90	1.7918 ^{***} (19.8)	1.6686 ^{***} (22.7)	1.7377 ^{***} (13.2)	1.7698 ^{***} (16.5)	1.5976 ^{***} (57.6)	1.5973 ^{***} (50.2)
DUR	.0226 ^{**} (16.3)	.0228 ^{**} (20.6)	.0143 ^{**} (8.38)	.0122 ^{**} (8.88)	.0031 ^{**} (11.1)	.0038 ^{**} (12.1)
RELQUAL	-1.0936 ^{**} (-3.18)	.8201 [*] (2.40)	-.2576 (-0.54)	.4789 (1.09)	1.8882 ^{**} (29.5)	-2.0627 ^{**} (-29.0)
UNCERT	5.5658 ^{***} (2.97)	3.4695 ^{**} (2.23)	6.0532 ^{**} (2.42)	2.7790 (1.38)	-	-
HOFF	.4669 ^{***} (4.13)	.2093 [*] (2.19)	.7671 ^{***} (4.65)	.2423 (1.54)	.5118 ^{***} (18.9)	-.4211 ^{***} (-8.86)
AOFF	.4039 ^{***} (3.06)	1.0484 ^{***} (12.2)	.6827 ^{***} (3.62)	1.4214 ^{***} (11.6)	-.3279 ^{***} (-6.41)	.5074 ^{***} (13.3)
DIFF-3	-.1087 (-0.48)	-1.0457 ^{***} (-3.58)	.2880 (1.07)	-.7491 ^{**} (-2.08)	.1076 [*] (1.79)	.0797 (1.34)
DIFF-2	.4881 ^{**} (4.63)	-.1867 (-1.62)	.6537 ^{***} (4.75)	-.2895 [*] (-1.70)	.1516 ^{***} (4.77)	-.0529 (-1.45)
DIFF-1	.3206 ^{***} (4.35)	-.0634 (-0.92)	.4186 ^{***} (4.39)	-.0968 (-1.03)	.0770 ^{***} (4.21)	-.0498 ^{**} (-2.34)
DIFF+1	-.0880 (-1.13)	.3469 ^{***} (6.20)	-.1548 (-1.47)	.4468 ^{***} (6.07)	-.0198 (-1.18)	.1167 ^{***} (6.16)
DIFF+2	-.2777 ^{**} (-2.35)	.2089 ^{***} (2.65)	-.0732 (-0.50)	.3253 ^{***} (3.06)	-.0265 (-1.07)	.1213 ^{***} (4.23)
DIFF+3	-.7561 ^{***} (-3.64)	.0197 (0.17)	-.8418 ^{***} (-2.82)	.2129 (1.35)	.0374 (1.08)	.0942 [*] (2.14)
Constant	-9.2644 ^{***} (-16.2)	-9.4477 ^{***} (-17.8)	-9.9234 ^{***} (-12.7)	-9.0710 ^{***} (-13.3)	-5.4499 ^{***} (-134.2)	-3.5043 ^{***} (-81.9)
Number of red cards/goals	1,323	2,073	786	1,177	26,903	20,219

Note: The hazard functions are estimated using in-play data on 18,324 English Premier League and Football League matches played during seasons 2001/02 to 2009/10 (inclusive). The timings of red cards and goals are recorded at frequencies of one minute. ^{***} denotes coefficient significantly different from zero, two-tail test, 1% level; ^{**} 5% level; ^{*} 10% level.

Table 2: Unconditional Distributions of the First- and Second-Half Differences Between the Home and Away Teams in Goals Scored and Red Cards Incurred

$y \rightarrow$	First-half difference in goals scored G(1)	Second-half difference in goals scored G(2)–G(1)	First-half difference in red cards incurred D(1)	Second-half difference in red cards incurred D(2)–D(1)
$y > 0$	0.344	0.370	0.017	0.043
$y = 0$	0.424	0.371	0.957	0.884
$y < 0$	0.232	0.258	0.026	0.073

Note: The data refer to 18,324 English Premier League and Football League matches played during seasons 2001/02 to 2009/10 (inclusive). For $y=G(1)$, the entry for $y>0$ (for example) reports the proportion of matches in which the home team scored more goals than the away team during the first half. For $y=G(2)–G(1)$, the entry for $y>0$ reports the proportion of matches in which the home team scored more goals than the away team during the second half. Similarly, for $y=D(1)$ and for $y=D(2)–D(1)$, the entries for $y>0$ report the proportions of matches in which the home team incurred more red cards than the away team in the first half and in the second half, respectively.

for the 45th or the 90th minute are several times the magnitudes of the hazards for other periods of one minute duration. Previous empirical evidence suggests that players from lower-quality teams and players from away teams are at greater risk of disciplinary sanction, and that the incidence of disciplinary sanction is higher in matches between evenly-balanced teams (Dawson et al, 2007). The coefficients on RELQUAL and UNCERT, a standard measure of uncertainty of match outcome, are signed accordingly and significant. The coefficients on HOFF and AOFF are positively signed and significant. The award of a red card increases the probabilities that either team incurs further red cards over the remaining duration of the present match.

The coefficients on DUR reflect a tendency, also evident in Table 2 above, for the arrival rate of red cards to increase over the match duration. As the match progresses, there is a tendency for cooperative (non-violent) behaviour to be superseded by non-cooperative (violent) behaviour. The coefficients on DIFF–2 and DIFF–1 are positive and significant in the home-team red card hazard function. The coefficient on DIFF+3 is insignificant. Likewise the coefficients on DIFF+1 and DIFF+2 in the away-team red card hazard are positive and significant, and the coefficient on DIFF–3 is insignificant. The probability of incurring a red card is higher for a team that is trailing in goals scored than it is for the same team when the scores are level.

Teams that are trailing are willing to take risks by switching to a violent style of play, and bearing an increased probability of incurring a red card, in order to increase the probability of scoring. Plausibly, however, this willingness to bear additional risk seems to be dependent on

the prospects of recovery to avoid defeat. Teams that are trailing by one or two goals perceive a realistic chance of recovery, and are more likely to take risks than teams trailing by larger margins with no realistic chance of recovery.

The coefficients on DIFF+1, DIFF+2 and DIFF+3 in the home-team red card hazard are negatively signed, and the latter two coefficients are significant. Likewise the coefficients on DIFF-1, DIFF-2 and DIFF-3 in the away-team red card hazard are negatively signed, and the coefficient on DIFF-3 is significant. The probability of incurring a red card appears to be lower for a team that is leading in goals scored than it is for the same team when the scores are level, and especially so for a team that is leading by a large margin. Teams when leading appear to play cautiously so as to minimize both the probability of conceding a goal, and the probability of incurring a red card.

In general, the coefficients of the hazard functions for red cards excluding those for two yellow cards follow a similar pattern to the coefficients of the hazards for all red cards. After excluding the two yellow-card offences, the coefficients on DUR become smaller in absolute magnitude, while the coefficients on DIFF-2 and DIFF-1 in the home-team red card hazard, and the coefficients on DIFF+1 and DIFF+2 in the away-team red card hazard, become larger. These variations appear consistent with the suggestion that part of the measured dependence of total red card on match duration is spurious, owing to the inclusion of red cards awarded for two yellow-card offences. The component of the red cards data that is dependent on the teams' strategic choices is identified more clearly in the estimations from which red cards awarded for two yellow-card offences are excluded.⁹

2.2 Goal-Scoring Hazard Functions

In the goal-scoring hazard functions, goal-scoring rates are expected to be relatively low in the 1st minute (the start of the match), the 46th minute (the resumption of the match following the half-time interval), and in any minute immediately following the scoring of a goal by either team. In each case, play starts or resumes from a kickoff at the halfway line, and the teams regroup into balanced defensive formations when play kicks off. Goals are unlikely to be scored during the first few seconds following kickoff. The coefficients on KICKOFF are signed accordingly, and significant. The coefficients on M45 and M90 are positively signed and significant, as before. The covariate RELQUAL controls for the direct relationship

⁹ The data set also allows the estimation of hazard functions for the arrivals of yellow cards, incurred for less serious acts of foul play. Yellow cards represent an alternative to red cards as an indicator of a violent style of play. The punishments for incurring yellow cards, however, are less immediate: as noted above, two yellow cards incurred by the same player in a single match result in the player's dismissal. Suspensions for yellow cards are triggered when the cumulative number of yellow cards from the start of the soccer season exceeds thresholds that vary over the duration of the season. The empirical hazard functions for yellow cards, not reported in Table 1, are generally similar to those for red cards.

between team quality and goal scoring. The coefficients on HOFF and AOFF are signed as expected, and significant.¹⁰

The coefficients on DUR reflect a tendency for the rates of goal scoring to increase over the duration of the match. This pattern is consistent with a tendency for cooperative forms of behaviour in the early stages of match to yield to non-cooperative behaviour in the latter stages.

The coefficients on DIFF-2 and DIFF-1 in the home-team goals hazard function are positive and significant, while the coefficient on DIFF-3 is positive but only borderline significant. The coefficients on DIFF+1, DIFF+2 and DIFF+3 in the away-team goals hazard function are positive and significant. These coefficients are consistent with a tendency for teams that are trailing in goals to seek to increase the probability of scoring, by either playing violently, or attacking, or both. In the theoretical model below, this condition ensures that the team that is trailing in goals scored increases its own goal-scoring probability by more than that of its opponent when it switches to (violent, attack).

The coefficients on the dummy variables for teams that are leading in goals scored (DIFF+1, DIFF+2 and DIFF+3 in the home-team goals hazard, and DIFF-1, DIFF-2 and DIFF-3 in the away-team goals hazard) are varied in sign, and are all smaller in absolute magnitude than the corresponding coefficients on the dummy variables for teams that are trailing in goals scored. After controlling for team quality, the goal-scoring probabilities of teams that are leading in goals scored are lower than those of teams that are trailing, but are not significantly different from those of teams that are level in goals scored. Any effect appears too small to register as statistically significant in the empirical model.

3 A Game-Theoretic Model of Strategic Behaviour

The evolution of strategic behaviour over the duration of a soccer match was first examined by Palomino et al. (2000) in an unpublished working paper.¹¹ Subsequently Dobson and Goddard (2010) use numerical simulations to examine the properties of a theoretical model in

¹⁰ The claim that is sometimes made by soccer pundits or commentators that a team playing with fewer players often tends to perform better than a team playing with a full complement of players is a popular myth, for which there is no statistical foundation (see also Ridder et al., 1994).

¹¹ Elsewhere the strategic choice of risk in dynamic competitive situations is analysed by Cabral (2002), in the context of an infinite-period R&D game in which two firms either imitate or differentiate themselves from their competitor by choosing between projects whose outcomes for either firm are positively correlated; and by Hvide (2002) and Hvide and Kristiansen (2003), in the context of a tournament in which the competitors can influence both the expectation and the variance of their performance through their choices of effort and risk-taking.

which the teams choose continuously, over the duration of the match, between defensive and attacking formations, and between a non-violent and a violent playing style.^{12 13}

In this section we develop a stylized two-period game-theoretic model of the strategic choices, and the rewards and costs deriving from these choices, for two opposing soccer teams. The soccer match comprises two discrete time intervals, each of 45 minutes' notional duration and known as the first and second halves, which correspond to the two periods in the theoretical model. The two halves are separated by a 15-minute interval, known as half-time; the end of the match is known as full-time. The duration of each half may be extended for a few minutes to compensate for losses playing time through stoppages (for injured players to receive treatment, or for other reasons).

For a two-period model, explicit solutions are available for the teams' optimal strategic choices; by contrast, explicit solutions are not feasible in the case where these choices vary continuously. In practice, and as reflected in the empirical results reported in Section 2, team managers can and do adjust continuously, primarily by means of shouted instructions from the touchline while the match is in play. The half-time interval provides a natural opportunity for a more measured or fundamental adjustment, in the privacy of the team's dressing room. Accordingly, while the two-period theoretical model is a simplification, developed at a higher level of abstraction than the empirical model, it bears sufficient resemblance to the reality to be capable of representing meaningfully the variation in the structure of rewards and deterrents associated with alternative strategic choices at different stages of the match.

3.1 Model Specification

At the start of each half, we assume that each team selects either a defensive or an attacking team formation, and either a non-violent or a violent style of play. For the home team, $i=1$ denotes (defend, non-violent); $i=2$ denotes (defend, violent); $i=3$ denotes (attack, non-violent);

¹² It is reasonable to ask whether these two dimensions of strategy could be subsumed into a single choice, for example between an attacking and defensive style of play, leading to a simpler theoretical model. In this case, simplicity would be achieved at the cost of a reduction in the model's ability to accommodate observed patterns in the timings of both goals and red cards. For example, the increase in the incidence of red cards between the first and second halves is proportionately much larger than the increase in the rate of goal scoring between the first and second halves.

¹³ Other soccer applications of game theory focus on the penalty kick. Chiappori et al. (2002) examine the selection of strategies by kickers and goalkeepers, and derive various mixed-strategy equilibria. Palacios-Huerta (2003) presents further empirical evidence to support the hypothesis that the behaviour of kickers and goalkeepers is consistent with a mixed-strategy equilibrium and random selection of strategies. Buzzacchi and Pedrini (2014) test the hypothesis that differently talented players randomize when taking kicks. For other sports, Walker and Wooders (2001) examine the choice of direction of first serves by leading professional tennis players. Banerjee, Swinnen and Weersink (2007) examine the impact on team strategies in the National Hockey League of a change in the league points scoring system affecting matches that were tied at the end of regular time.

and $i=4$ denotes (attack, violent). For the away team, $j=1, \dots, 4$ denotes the same four choices. I_v and J_v are 0-1 indicator variables for the selection of violent play by the home and away teams, respectively; and I_a and J_a are indicator variables for the selection of attacking play.

The state of the match at any stage is defined by the current difference in goals scored, and the numbers of players from either team (if any) that have incurred red cards.¹⁴ In order to make the model tractable, we simplify further by assuming that the *only* feasible transitions in the state of the match during each half are (i) a change in the difference between the home-team and away-team scores of no more than one goal per half; and (ii) the award of a maximum of one red card per half. Although this theoretical set-up is highly stylized,¹⁵ it provides a representation of the reality that is adequate for the purpose of identifying the key strategic choices of the opposing teams at different stages of the match. The state of the match at the start (denoted $t=0$), at half-time ($t=1$) and at full-time ($t=2$) is defined by $G(t)$ and $D(t)$, where $G(t)$ is the difference between the teams in goals scored at stage t , $D_H(t)$ and $D_A(t)$ are the numbers of home- and away-team players who have incurred red cards, respectively, and $D(t)=D_H(t)-D_A(t)$.¹⁶

In the theoretical model, the transition probabilities for a change in the difference between the home- and away-team scores during the second half are dependent on the teams' choices of i and j for the second half, and on $D(1)$. The transition probabilities for a change in the difference between the home- and away-team scores during the first half are dependent only on the teams' choices of i and j for the first half, and are identical to the corresponding second-half transition probabilities for the case $D(1)=0$.

Let $p_{i,j,D(1)}$ denote the probability of a one-goal change in the difference in scores in favour of the home team during the second half, conditional on i , j and $D(1)$; and let $q_{i,j,D(1)}$ denote the probability of a one-goal change in the difference in scores in favour of the away team during the second half. Similarly $p_{i,j,0}$ and $q_{i,j,0}$ denote the probabilities of a one-goal change in the difference in scores in favour of the home and away teams during the first half.

¹⁴ A yellow card, also known as a booking or caution, is awarded when a player commits less serious forms of foul play. There is no further punishment within the match, unless the player commits a second offence, in which case a red card is awarded and the player is dismissed for the rest of the match (with no replacement permitted). A red card or dismissal is awarded for more serious offences, and results in immediate expulsion (again, with no replacement permitted). After the match, a red card leads to a suspension, preventing the player from appearing in one, two or three of his team's next scheduled matches.

¹⁵ In practice, the difference in scores might change by more than one goal per half, and more than one player might incur a red card. Furthermore, teams may select from a continuum of offensive and defensive strategies; and this selection may be varied continuously throughout the duration of the match.

¹⁶ The range of feasible values for $G(t)$ and $D(t)$ is as follows: $G(0), D(0)=0$, $G(1), D(1)={+1, 0, -1}$ and $G(2), D(2)={+2, +1, 0, -1, -2}$.

Let u_i denote the probability that a home-team player incurs a red card within a 45-minute period of play, and let v_j denote the corresponding away-team probability. In the theoretical model, u_i and v_j are assumed to be the same for both halves of the match, and dependent only on the teams' own choices of i or j for that half. The assumed specifications for $p_{i,j,D(t)}$, $q_{i,j,D(t)}$ (for $D(t) \in \{-1, 0, 1\}$), u_i and v_j are

$$p_{i,j,0} = s_H + b(I_a + 2J_a) + c(2I_v + J_v); \quad q_{i,j,0} = s_A + b(2I_a + J_a) + c(I_v + 2J_v) \quad (1)$$

$$p_{i,j,1} = p_{i,j,0} - d; \quad p_{i,j,-1} = p_{i,j,0} + d; \quad q_{i,j,1} = q_{i,j,0} + d; \quad q_{i,j,-1} = q_{i,j,0} - d \quad (2)$$

$$u_i = r_H + eI_a + fI_v; \quad v_j = r_A + eJ_a + fJ_v \quad (3)$$

All parameters in (1), (2) and (3) are positive, and are constrained so that all probabilities lie between zero and one. In [1], s_H and s_A are the baseline probabilities of a one-goal change in the difference in scores in favour of the home and away teams, respectively, for the cases $i,j=1$ (defend, non-violent), during 45 minutes' play. $s_H > s_A$ is consistent with the existence of a home-field advantage. By switching from defence to attack, the home team increases the probability of a one-goal change in the score in its favour by b , but it also increases the probability of a one-goal change in its opponent's favour by $2b$. The corresponding effects for the away team are the same.

The teams' strategic choices are challenging only if they entail trade-offs between reward and deterrent. We therefore assume that adopting an attacking style increases the probability of conceding a goal (deterrent) by more than it increases the probability of scoring (reward). This seems the most natural way of capturing the tendency of many teams to switch from a defensive to an offensive strategy after falling behind: the switch occurs because the team that is trailing has nothing left to lose (the deterrent becomes ineffective), and accepts a greater risk of falling further behind as an acceptable price to pay for the reward of an increase in the probability of drawing level. If, by attacking, a team could increase its own scoring probability by more than its opponent's, then presumably most teams would attack from the start of every match.

By switching from non-violence to violence in (1), the home team increases the probability of a one-goal change in the score in its favour by $2c$, and it increases the probability of a one-goal change in its opponent's favour by c . The corresponding effects for the away team are the same. Accordingly, adopting a violent style, or cheating by playing outside the rules of the sport, is assumed to increase the probability of scoring a goal (reward) by more than it increases the probability of conceding (part of the deterrent), so long as this behaviour goes unpunished. There is, however, an additional deterrent in the event of punishment in the form of a red card, through an increased probability of conceding goals

subsequently in the current match, which recedes over the match duration, as well as a weakened team selection in future matches.

In (2), a home team that starts the second half at a numerical disadvantage in players incurs a reduction of d in the probability of a one-goal change in its favour. The probability of a one-goal change in its opponent's favour is increased by d . The corresponding effects for the away team are the same.

In (3), r_H and r_A are the baseline probabilities of incurring a red card for a home-team player and an away-team player, respectively, for the cases $I_j=1$ (defend, non-violent), during 45 minutes' play. $R_A > r_H$ is consistent with an observed tendency for away team players to be 299 penalized more frequently than home team players, the teams' strategic choices notwithstanding. By switching from defence to attack, the home team increases its own probability of incurring a red card by e ; there is no effect on the corresponding probability for the opposing team. By switching from non-violence to violence, the home team increases its own probability of incurring a red card by f ; again, there is no effect for the opponent. The corresponding effects for the away team are the same.

The payoff for each team at the end of the match is dependent on the league points gained from the match, and the cost of future player suspensions arising from red cards incurred during the match (if any). The payoffs from the league points gained are determined through a utility function, which may be either linear in points gained (risk-neutrality) or concave (risk-aversion). The available league points are 3 for a win, 1 for a draw (tie), and 0 for a loss. The corresponding utility payoffs are m , 1 and 0, respectively, where $m=3$ represents risk neutrality and $1 < m < 3$ represents risk aversion. The utility cost of red cards (arising from future league points lost due to player suspensions) is a fixed utility deduction per red card incurred, denoted $n > 0$.

The home- and away-team full-time payoffs, denoted $\bar{h}_{(2)}[G(2), D_H(2)]$ and $\bar{a}_{(2)}[G(2), D_A(2)]$, are the summation of the utility payoff from the league points gained from the present match, and the utility cost of any red cards incurred in the present match:

$$\begin{aligned} \bar{h}_{(2)}[G(2), D_H(2)] &= m - nD_H(2); \bar{a}_{(2)}[G(2), D_A(2)] = -nD_A(2) \text{ for } G(2)=2, 1 \\ \bar{h}_{(2)}[G(2), D_H(2)] &= 1 - nD_H(2); \bar{a}_{(2)}[G(2), D_A(2)] = 1 - nD_A(2) \text{ for } G(2)=0 \\ \bar{h}_{(2)}[G(2), D_H(2)] &= -nD_H(2); \bar{a}_{(2)}[G(2), D_A(2)] = m - nD_A(2) \text{ for } G(2)=-1, -2 \end{aligned} \quad (4)$$

3.2 Calibration of Parameters

Tables 2, 3 and 4 report some summary statistics based on the data set used to estimate the empirical model reported in Section 2, configured for consistency with the theoretical model described above. Table 2 reports the unconditional empirical distributions of the first- and

second-half differences in goals scored and red cards. The probabilities that either team scores more goals than its opponent in the second half are marginally higher than the corresponding first-half probabilities. The probabilities that either team incurs more red cards than its opponent in the second half are substantially higher than the first-half probabilities.

Table 3 reports the empirical distributions of second-half differences in goals scored, conditional on the first-half differences in goals scored and red cards incurred. The probabilities that either team scores more second-half goals than its opponent conditional on trailing in goals scored at half-time are larger than the corresponding probabilities conditional on leading at half-time. This is consistent with the notion that a team that is trailing takes risks in order to improve its scoring rate, while a team that is leading avoids risk.¹⁷ The probabilities that either team scores more second-half goals than its opponent, conditional on being at a numerical advantage or disadvantage in players at half-time, demonstrate that a team incurring a red card during the first half typically bears a heavy cost, in the form of a substantial increase in the probability of defeat.

Table 4 reports the empirical distributions of second-half differences in red cards incurred, conditional on the half-time differences in goals scored and red cards incurred. The probabilities that either team incurs more second-half red cards than its opponent conditional on trailing in goals scored at half-time are larger than the corresponding probabilities conditional on leading at half-time. As before, this pattern is consistent with the hypothesis that teams tend to take risks when trailing in goals scored, and avoid risk when leading.

In accordance with the data reported in Tables 2, 3 and 4, and solely for the purpose of developing the numerical illustrations that are presented in the next section, we posit the following ‘ballpark’ parameter estimates in (1)-(4):

$$\begin{aligned} s_H=0.34, s_A=0.23, r_H=0.02, r_A=0.03, b=0.01, \\ c=0.02, d=0.12, e=0.015, f=0.03, m=2, n=0.5 \end{aligned} \quad (5)$$

The parameters r_H and r_A are calibrated in accordance with the empirical model, which indicates that after controlling for goal difference in the current match home teams incur a lower rate of player dismissals than away teams (see also Dawson et al., 2007).¹⁸

¹⁷ These data, however, do not provide an unambiguous guide to the calibration of the parameters b and c in (1). Matches in which the home team is leading at half-time (for example) contain a higher-than-average proportion of cases where the home team is stronger than the away team. This suggests that the differences between the scoring probabilities of teams taking risks and avoiding risks are higher than the differences between the empirical second-half scoring rates of teams that are leading or trailing in goals scored at half-time.

¹⁸ The probabilities that either team incurs more second-half red cards than its opponent conditional on being at a numerical advantage in players at half-time are larger than the corresponding probabilities conditional on being at a numerical disadvantage. This might reflect a tendency for teams that have

Table 3: Conditional Distributions of the Second-Half Difference Between the Home and Away Teams in Goals Scored

$y \rightarrow$ $x \rightarrow$	Second-half difference in goals scored, $G(2)-G(1)$					
	$G(1)>0$	$G(1)=0$	$G(1)<0$	$D(1)>0$	$D(1)=0$	$D(1)<0$
$y > 0 x$	0.372	0.358	0.391	0.248	0.368	0.520
$y = 0 x$	0.359	0.390	0.357	0.363	0.373	0.329
$y < 0 x$	0.269	0.252	0.254	0.389	0.259	0.152
No. of matches	6,310	7,765	4,249	303	17,540	481

Note: The data refer to 18,324 English Premier League and Football League matches played during seasons 2001/02 to 2009/10 (inclusive). For $x=G(1)>0$, the entry for $y > 0 | x$ (for example) reports, of those matches in which the home was leading in goals scored at half-time, the proportion in which the home team scored more goals than the away team in the second half. For $x=G(1)=0$, the entry for $y > 0 | x$ reports, of those matches in which the scores were level at half-time, the proportion in which the home team scored more goals than the away team in the second half. For $x=G(1)<0$, the entry for $y > 0 | x$ reports, of those matches in which the away team was leading at half-time, the proportion in which the home team scored more goals than the away team in the second half. Similarly for $x=D(1)>0, =0$ or <0 , the entries for $y > 0 | x$ report the proportions of matches in which the home team scored more goals than the away team in the second half conditional on the difference between the teams in red cards incurred during the first half.

The justification for the specified values for the parameters m and n is as follows. By setting $m=2$, we assume that the unobservable difference in utility between a win and a draw is the same as the difference between a draw and a loss. Given that the difference in league points awarded for a win and a draw is twice the difference between a draw and a loss, this is tantamount to an assumption of risk-aversion on the part of team managers.¹⁹ The utility cost

been on the receiving end of violent play by an opponent to attempt to exact retribution; or a tendency for referees who have already dismissed a player from one team to look for the opportunity to restore numerical equality by dismissing a player from the opposing team.

¹⁹The observation that teams often tend to play defensively when the scores are level in an effort to preserve one league point, rather than risk losing a point by playing more offensively in an effort to gain all three points, suggests a concave utility function is a reasonable conjecture. Observation of numerous post-match TV interviews given by managers and players suggests avoidance of defeat is often a high priority, of similar importance to the conversion of draws into wins. The award of three, rather than two, league points for a win, was introduced in most countries during the 1980s and 1990s in an effort to curb a tendency for many teams to play defensively most of the time, by increasing the incentive to attack. The perceived need for such a measure is, in itself, circumstantial evidence of an attitude of risk aversion on the part of managers and coaches. In the English Premier League, an average return of one league point per match is usually just sufficient to ensure an end-of-season standing above the bottom three places, thereby avoiding relegation to the second tier and huge financial penalties in lost TV and other revenue. The threat of a collapse in form and a run of consecutive defeats leading to relegation hangs over most teams for most of the season; it is usually

Table 4: Conditional Distributions of 2nd Half Difference in Dismissals

$y \rightarrow$ $x \rightarrow$	2nd half difference in dismissals, $D(2)-D(1)$					
	$G(1)>0$	$G(1)=0$	$G(1)<0$	$D(1)>0$	$D(1)=0$	$D(1)<0$
$y > 0 x$	0.033	0.043	0.059	0.030	0.043	0.064
$y = 0 x$	0.883	0.884	0.884	0.865	0.885	0.867
$y < 0 x$	0.084	0.073	0.057	0.106	0.072	0.069
No. of matches	6,310	7,765	4,249	303	17,540	481

Note: The data refer to 18,324 English Premier League and Football League matches played during seasons 2001/02 to 2009/10 (inclusive). For $x=G(1)>0$, the entry for $y > 0 | x$ (for example) reports, of those matches in which the home was leading in goals scored at half-time, the proportion in which the home team incurred more red cards than the away team in the second half. For $x=G(1)=0$, the entry for $y > 0 | x$ reports, of those matches in which the scores were level at half-time, the proportion in which the home team incurred more red cards than the away team in the second half. For $x=G(1)<0$, the entry for $y > 0 | x$ reports, of those matches in which the away team was leading at half-time, the proportion in which the home team incurred more red cards than the away team in the second half. Similarly for $x=D(1)>0, =0$ or <0 , the entries for $y > 0 | x$ report the proportions of matches in which the home team incurred more red cards than the away team in the second half conditional on the difference between the teams in red cards incurred during the first half.

of player suspensions arising from red cards is also unobservable. The league points or utility cost of player suspensions arising from red cards depends on the size and quality of the squad of players from which the team is selected, the quality of the suspended player, and the number of matches from which the player is suspended (usually either one or three matches). We assume that each suspension causes an expected loss of 0.5 league points.

3.3 Derivation of the Dominant-Strategy Equilibria

As is standard practice, the derivation of the solutions for the teams' optimal strategies in the second- and first-half sub-games proceeds via backward iteration. Below, we specify the conditions on the model's parameter values that lead to the emergence of a Pareto suboptimal equilibrium with a prisoner's dilemma structure in the second half of a match that is level in goals scored at half-time, preceded by a Pareto optimal equilibrium in the first half.

only at a relatively late stage that the majority of teams cross the league points threshold required to ensure a finishing position above the relegation places (in the Premier League, normally 35-40 points from 38 matches). While it may seem that three league points gained from one win and two defeats should substitute perfectly for three points gained from three draws, the fear of relegation appears to instill a defensive psychology that replicates the concave payoff structure of the end-of-season league standings into the utility value of the league points gained from each match.

For the second-half sub-game, the following expressions govern the relations between the expectations at half time of the teams' full-time payoffs conditional on the values of i and j selected for the second half, denoted $h_{(1)}^{i,j}[G(1), D_H(1), D_A(1)]$ and $a_{(1)}^{i,j}[G(1), D_H(1), D_A(1)]$, and the teams' full-time payoffs unconditional on i and j :

$$h_{(1)}^{i,j}[G(1), D_H(1), D_A(1)] = \mathbf{\Gamma}' \mathbf{\Omega}_{(2)} \mathbf{\Lambda}; \quad a_{(1)}^{i,j}[G(1), D_H(1), D_A(1)] = \mathbf{\Gamma}' \mathbf{\Phi}_{(2)} \mathbf{\Lambda}$$

where $\mathbf{\Gamma}' = (q_{i,j,D(1)} \quad 1 - p_{i,j,D(1)} - q_{i,j,D(1)} \quad p_{i,j,D(1)})$; $\mathbf{\Lambda}' = (v_j \quad 1 - u_i - v_j \quad u_i)$

$$\mathbf{\Omega}_{(2)} = \begin{pmatrix} \bar{h}_{(2)}[G(1) - 1, D_H(1)] & \bar{h}_{(2)}[G(1) - 1, D_H(1)] & \bar{h}_{(2)}[G(1) - 1, D_H(1) + 1] \\ \bar{h}_{(2)}[G(1), D_H(1)] & \bar{h}_{(2)}[G(1), D_H(1)] & \bar{h}_{(2)}[G(1), D_H(1) + 1] \\ \bar{h}_{(2)}[G(1) + 1, D_H(1)] & \bar{h}_{(2)}[G(1) + 1, D_H(1)] & \bar{h}_{(2)}[G(1) + 1, D_H(1) + 1] \end{pmatrix}$$

$$\mathbf{\Phi}_{(2)} = \begin{pmatrix} \bar{a}_{(2)}[G(1) - 1, D_A(1) + 1] & \bar{a}_{(2)}[G(1) - 1, D_A(1)] & \bar{a}_{(2)}[G(1) - 1, D_A(1)] \\ \bar{a}_{(2)}[G(1), D_A(1) + 1] & \bar{a}_{(2)}[G(1), D_A(1)] & \bar{a}_{(2)}[G(1), D_A(1)] \\ \bar{a}_{(2)}[G(1) + 1, D_A(1) + 1] & \bar{a}_{(2)}[G(1) + 1, D_A(1)] & \bar{a}_{(2)}[G(1) + 1, D_A(1)] \end{pmatrix} \quad (6)$$

In the case $G(1)=0, D_H(1)=D_A(1)=0$, [6] leads to the following:²⁰

$$\begin{aligned} h_{(1)}^{1,1} &= 1 + (m - 1)s_H - s_A - nr_H; & h_{(1)}^{1,2} &= h_{(1)}^{1,1} + cm - 3c; \\ h_{(1)}^{1,3} &= h_{(1)}^{1,1} + 2bm - 3b; & h_{(1)}^{1,4} &= h_{(1)}^{1,2} + h_{(1)}^{1,3} - h_{(1)}^{1,1}; \\ h_{(1)}^{2,j} &= h_{(1)}^{1,j} + 2cm - 3c - fn; & h_{(1)}^{3,j} &= h_{(1)}^{1,j} + bm - 3b - en; \\ h_{(1)}^{4,j} &= h_{(1)}^{2,j} + h_{(1)}^{3,j} - h_{(1)}^{1,j} \text{ for } j=1, \dots, 4 \end{aligned} \quad (7)$$

For the away team, $a_{(1)}^{i,j}$ is obtained by replacing s_A, s_H, r_A and r_H with s_H, s_A, r_H and r_A , respectively, in the expression for $h_{(1)}^{j,i}$ in [7].

Accordingly, $\{i=2, j=2\}$ is a dominant strategy equilibrium and Nash equilibrium under the following conditions:

$$\begin{aligned} c(2m-3) - fn &> 0; \\ -3b(m-1) + c(2m-3) + (e-f)n &> 0; \\ b(3-m) + c(2m-3) + en &> 0 \end{aligned} \quad (8)$$

²⁰ The arguments of the conditional payoff functions reported below are omitted for convenience.

This dominant strategy equilibrium, at which both teams select (defend, violent), has a prisoner’s dilemma structure and is Pareto suboptimal if $h_{(1)}^{1,1} > h_{(1)}^{2,2}$ and $a_{(1)}^{1,1} > a_{(1)}^{2,2}$. A necessary and sufficient condition is

$$3c(m-2) - fn < 0 \tag{9}$$

Conditions (8) and (9) are satisfied by the parameter values in (5). For illustrative purposes, Table 5 reports numerical evaluations of the home- and away-team payoffs for the second-half sub-game based on (5). Both teams would be better off in comparison with the dominant strategy equilibrium $\{i=2,j=2\}$ if both were to select (defend, non-violent), leading to the cooperative solution $\{i=1,j=1\}$. The cooperative solution is unstable, however, because both teams would be able to improve their payoffs further through defection, from non-violence to violence, provided the other team does not simultaneously defect.

Conditions (8) and (9) ensure the existence of a Pareto-suboptimal dominant strategy equilibrium $\{i=2,j=2\}$ in the cases where $G(1)=0$ and either $D_H(1)=1$ or $D_A(1)=1$.²¹ In the case $G(1)=-1$ and for any feasible $D_H(1)$ and $D_A(1)$, $\{i=4,j=1\}$ is a dominant-strategy equilibrium under any set of feasible parameter values in (5). The payoffs in the case $G(1)=-1$, $D_H(1)=D_A(1)=0$ are reported in the Appendix. By symmetry, $\{i=1,j=4\}$ is a dominant-strategy equilibrium in the case $G(1)=1$. The team that is trailing in goals scored at half time selects

Table 5: Matrix of Expected Payoffs at Half-Time Conditional on Second-Half Strategies: Scores Level, No Red Cards [$G(1)=0$, $D_H(1)=D_A(1)=0$]

	Expected payoffs							
	Home team = $h_{(i)}^{i,j}(0,0,0)$				Away team = $a_{(i)}^{i,j}(0,0,0)$			
	j = 1	j = 2	j = 3	j = 4	j = 1	j = 2	j = 3	j = 4
i = 1	1.1000	1.0800	1.1100	1.0900	0.8750	0.8800	0.8575	0.8625
2	1.1050	1.0850	1.1150	1.0950	0.8550	0.8600	0.8375	0.8425
3	1.0825	1.0625	1.0925	1.0725	0.8850	0.8900	0.8675	0.8725
4	1.0875	1.0675	1.0975	1.0775	0.8650	0.8700	0.8475	0.8525

Note: Payoffs for dominant strategies are highlighted in **bold**. Dominant strategy equilibrium is highlighted in **bold italics**. Strategies are as follows: $i,j=1$ denotes (defend, non-violent); $i,j=2$ denotes (defend, violent), $i,j=3$ denotes (attack, non-violent), $i,j=4$ denotes (attack, violent). The parameter values in [1]-[4] are defined by [5].

²¹ According to (2), the award of a red card in the first half produces the same decrease in his team’s goal-scoring rate in the second half regardless of its choice of strategy, and the same increase in the opponent’s goal-scoring rate regardless of strategy. The values of i and j that define the equilibrium in the second-half sub-game are therefore unaffected by the values of $D_H(1)$ and $D_A(1)$.

Table 6: Matrix of Expected Payoffs at Half-Time Conditional on Second-Half Strategies: Away Team Leads By One Goal, No Red Cards [$G(1)=-1, D_H(1)=D_A(1)=0$]

		Expected payoffs							
		Home team = $h_{(1)}^{i,j}(-1,0,0)$				Away team = $a_{(1)}^{i,j}(-1,0,0)$			
		$j=1$	$j=2$	$j=3$	$j=4$	$j=1$	$j=2$	$j=3$	$j=4$
$i=1$.3300	.3500	.3500	.3700	1.6450	1.6100	1.6175	1.5825
	2	.3550	.3750	.3750	.3950	1.6050	1.5700	1.5775	1.5425
	3	.3325	.3525	.3525	.3725	1.6350	1.6000	1.6075	1.5725
	4	.3575	.3775	.3775	.3975	1.5950	1.5600	1.5675	1.5325

Notes: See Table 5.

Table 7: Matrix of Expected Payoffs at Half-Time Conditional on Second-Half Strategies: Home Team Leads By One Goal, No Red Cards [$G(1)=1, D_H(1)=D_A(1)=0$]

		Expected payoffs							
		Home team = $h_{(1)}^{i,j}(1,0,0)$				Away team = $a_{(1)}^{i,j}(1,0,0)$			
		$j=1$	$j=2$	$j=3$	$j=4$	$j=1$	$j=2$	$j=3$	$j=4$
$i=1$		1.7600	1.7200	1.7500	1.7100	0.2150	0.2400	0.2175	0.2425
	2	1.7250	1.6850	1.7150	1.6750	0.2350	0.2600	0.2375	0.2625
	3	1.7325	1.6925	1.7725	1.6825	0.2350	0.2600	0.3375	0.3625
	4	1.6975	1.6575	1.6875	1.6475	0.2550	0.2800	0.2575	0.2825

Notes: See Table 5.

(attack, violent) in the second half, while the team that is leading selects (defend, non-violent). Based on the illustrative parameter values defined by (5), Tables 6 and 7 report the home- and away-team payoffs for the second-half sub-games in the cases $G(1)=-1, D_H(1)=D_A(1)=0$; and $G(1)=1, D_H(1)=D_A(1)=0$.

In the first-half sub-game, the expectations at half time of the full-time payoffs, unconditional on i and j , are

$$\bar{h}_{(1)}[G(1), D_H(1), D_A(1)] = \max_i h_{(1)}^{i,j^*}[G(1), D_H(1), D_A(1)] \quad (10)$$

$$\bar{a}_{(1)}[G(1), D_H(1), D_A(1)] = \max_j a_{(1)}^{i^*,j}[G(1), D_H(1), D_A(1)]$$

where $i^* = \arg \max_i h_{(1)}^{i,j^*}[G(1), D_H(1), D_A(1)]$

and $j^* = \arg \max_j a_{(1)}^{i^*,j}[G(1), D_H(1), D_A(1)]$

In the first-half sub-game, the following expressions define the expected payoffs at the start of the match *conditional* on the teams' choices of i and j for the first half, denoted $h_{(0)}^{i,j}$ and $a_{(0)}^{i,j}$

$$h_{(0)}^{i,j} = \Gamma_0' \Omega_{(1)} \Lambda; \quad a_{(0)}^{i,j} = \Gamma_0' \Phi_{(1)} \Lambda$$

where $\Gamma_0' = \Gamma' = (q_{i,j,0} \quad 1-p_{i,j,0}-q_{i,j,0} \quad p_{i,j,0})$, Λ is defined as in [6], and

$$\begin{aligned} \Omega_{(1)} &= \begin{pmatrix} \bar{h}_{(1)}[-1,0,1] & \bar{h}_{(1)}[-1,0,0] & \bar{h}_{(1)}[-1,1,0] \\ \bar{h}_{(1)}[0,0,1] & \bar{h}_{(1)}[0,0,0] & \bar{h}_{(1)}[0,1,0] \\ \bar{h}_{(1)}[1,0,1] & \bar{h}_{(1)}[1,0,0] & \bar{h}_{(1)}[1,1,0] \end{pmatrix} \\ \Phi_{(1)} &= \begin{pmatrix} \bar{a}_{(1)}[-1,0,1] & \bar{a}_{(1)}[-1,0,0] & \bar{a}_{(1)}[-1,1,0] \\ \bar{a}_{(1)}[0,0,1] & \bar{a}_{(1)}[0,0,0] & \bar{a}_{(1)}[0,1,0] \\ \bar{a}_{(1)}[1,0,1] & \bar{a}_{(1)}[1,0,0] & \bar{a}_{(1)}[1,1,0] \end{pmatrix} \end{aligned} \quad (11)$$

The payoffs derived from (11) are reported in the Appendix. Though the algebraic expressions for the payoffs are lengthy, the derivation of the following conditions, under which $i=1$ is the home team's dominant strategy, is straightforward:

$$\begin{aligned} fz_H + c[2m - ne + 2nf - 3 + s_H(4-3m) + b(3-2m) + c(24-13m) \\ + s_A(5-2m) + d(r_H-r_A+f)(1+m)] < 0 \\ ez_H + b[m - 2ne + nf - 3 + s_H(5-3m) + b(3-m) + c(24-11m) \\ + s_A(4-m) + d(r_H-r_A+f)(2m-1)] < 0 \end{aligned} \quad (12)$$

where $z_H = ds_H - ds_A + dms_A - dm - n$

By symmetry, the corresponding conditions under which $j=1$ is the away team's dominant strategy are obtained by replacing s_A, s_H, r_A and r_H with s_H, s_A, r_H , and r_A , respectively, and z_H with z_A (similarly defined), in (12).

Condition (12), and the corresponding condition for the away team, is satisfied by the parameter values in (5). The dominant strategy for both teams in the first half is (defend, non-violent). $\{i=1, j=1\}$ is a dominant strategy equilibrium, and a Nash equilibrium. This cooperative solution is Pareto-optimal. Based on the illustrative parameter values defined by (5), Table 8 reports the home- and away-team payoffs for the first-half sub-game.

The game-theoretic model yields a number of propositions consistent with patterns in the data as described in Section 2. In particular, (i) the propensity for risk-taking behaviour leading to goals being scored or players being dismissed increases with the duration of the

Table 8: Matrix of Expected Payoffs at Start of Match Conditional on First-Half Strategies

Expected payoffs at start of match, conditional on first-half strategies								
	Home team = $h_{(0)}^{i,j}$				Away team = $a_{(0)}^{i,j}$			
	$j = 1$	$j = 2$	$j = 3$	$j = 4$	$j = 1$	$j = 2$	$j = 3$	$j = 4$
$i = 1$	1.1219	1.1102	1.1296	1.1176	0.8024	0.7996	0.7874	0.7848
2	1.1123	1.1006	1.1200	1.1081	0.7974	0.7946	0.7824	0.7798
3	1.1035	1.0918	1.1112	1.0993	0.8135	0.8107	0.7985	0.7959
4	1.0942	1.0825	1.1019	1.0900	0.8083	0.8054	0.7933	0.7907

Notes: See Table 5.

match, which follows from the decrease in the cost of incurring a red card over the duration of the match.

Accordingly, the incentive to attack or play violently (or both) increases with the duration; (ii) after controlling for team-quality effects, the propensity to incur a red card is greater for teams that are trailing in goals scored than it is for teams that are level or leading. This reflects the anticipated effect of attacking or violent play on the part of a team that is trailing in goals scored on the team’s own probability of incurring a red card; (iii) after controlling for team-quality effects, the propensity to score a goal is greater for teams that are trailing in goals scored than it is for teams that are level or leading. This reflects the effect on a team’s goal-scoring rate of an increased propensity to attack and to play violently when the team is trailing in goals scored. If the team that is trailing in goals scored selects (violent, attack) and the team that is leading selects (non-violent, defend), (iii) follows if $c > b$ in (5), as we have assumed in the illustrations presented in the previous section.

4 Conclusion

This study presents novel non-experimental empirical and theoretical evidence on one of the fundamental propositions of game theory that cooperative behaviour tends to break down when the configuration of rewards and costs to the players are such as to create payoffs with a prisoner’s dilemma structure. While the literature presents extensive experimental evidence on the adoption of cooperative and non-cooperative forms of behaviour, the results of this study augment a relatively sparse extant body of non-experimental evidence.

We present empirical evidence on the timings of red cards and goals from a sample of more than 18,000 professional soccer matches in the English leagues. After controlling for team-quality effects, goal-scoring rates for teams that are currently trailing by one or two goals are significantly higher than the rates for teams that are leading or level in goals scored. Teams that are currently trailing are willing to bear an increased probability of incurring a red card, in order to increase the probability of scoring a goal quickly. By contrast, teams that are

currently leading tend to play defensively and non-violently. We interpret a pronounced increase in the propensity for violent play to take place in the latter stages of soccer matches, evidenced by a marked increase in the incidence of red cards, as novel non-experimental evidence in support of one of the key propositions of game theory, that players typically resort to mutually detrimental non-cooperative forms of behaviour when the payoffs assume a prisoner's dilemma structure.

We develop a stylized two-period game-theoretic model of the strategic choices made by soccer teams when selecting between defensive and attacking team formations, and between non-violent and violent styles of play. The two periods correspond to the first and second halves of a regular 90 minute soccer match. According to the theoretical analysis, optimal strategic behaviour in the second half of the match depends on the difference in scores at half-time. Teams that are trailing in goals scored tend to attack and play violently in the second half. When the half-time scores are level, under reasonable assumptions the second-half sub-game assumes a prisoner's dilemma structure, and there is a tendency for both teams to defect from cooperative (non-violent) behaviour in the first half to non-cooperative (violent) behaviour in the second half. This feature of the model is consistent with an observed and hitherto unexplained tendency for the incidence of red cards to increase markedly during the latter stages of matches.

Appendix

This appendix reports the payoffs for several of the sub-games discussed in Section 3.

A. *Second-half sub-game*: $G(1)=-1$, $D_H(1)=D_A(1)=0$

(6) leads to the following payoffs:

$$\begin{aligned} h_{(1)}^{1,1} &= s_H - nr_H; \quad h_{(1)}^{1,2} = h_{(1)}^{1,1} + c; \quad h_{(1)}^{1,3} = h_{(1)}^{1,1} + 2b; \quad h_{(1)}^{1,4} = h_{(1)}^{1,2} + h_{(1)}^{1,3} - h_{(1)}^{1,1} \\ h_{(1)}^{2,j} &= h_{(1)}^{1,j} + 2c - nf; \quad h_{(1)}^{3,j} = h_{(1)}^{1,j} + b - ne; \quad h_{(1)}^{4,j} = h_{(1)}^{2,j} + h_{(1)}^{3,j} - h_{(1)}^{1,j} \quad \text{for } j=1,\dots,4 \\ a_{(1)}^{1,1} &= m - nr_A - (m-1)s_H; \quad a_{(1)}^{2,1} = a_{(1)}^{1,1} - 2c(m-1); \quad a_{(1)}^{3,1} = a_{(1)}^{1,1} - b(m-1); \\ a_{(1)}^{4,1} &= a_{(1)}^{2,1} + a_{(1)}^{3,1} - a_{(1)}^{1,1}; \\ a_{(1)}^{i,2} &= a_{(1)}^{i,1} - fn - c(m-1); \quad a_{(1)}^{i,3} = a_{(1)}^{i,1} - en - 2b(m-1); \quad a_{(1)}^{i,4} = a_{(1)}^{i,2} + a_{(1)}^{i,3} - a_{(1)}^{i,1} \\ &\quad \text{for } i=1,\dots,4 \end{aligned}$$

It follows that $i=4$ (attack, violent) is the home team's dominant strategy under the conditions $2c - nf > 0$ and $b - ne > 0$; and $j=1$ (defend, non-violent) is the away team's dominant strategy

under any feasible permutation of parameter values. Accordingly, $\{i=4, j=1\}$ is a dominant-strategy equilibrium.

B. Second-half sub-game: $G(1)=1, D_H(1)=D_A(1)=0$

(6) leads to the following payoffs: $h_{(1)}^{i,j}$ is defined by replacing s_A, s_H, r_A and r_H with s_H, s_A, r_H and r_A , respectively, in $a_{(1)}^{j,i}$ in the second-half sub-game for $G(1)=-1, D_H(1)=D_A(1)=0$; and $a_{(1)}^{i,j}$ is similarly defined with reference to $h_{(1)}^{j,i}$. By symmetry, $\{i=1, j=4\}$ is a dominant-strategy equilibrium under the conditions $2c - nf > 0$ and $b - ne > 0$.

C. First-half sub-game

With reference to (10), and in accordance with the preceding analysis of the second-half sub-games, we assume $\{i^*, j^*\} = \{4, 1\}$ for $G(1)=-1, \{2, 2\}$ for $G(1)=0$ and $\{1, 4\}$ for $G(1)=1$.

Substitution of $\{i^*, j^*\}$ into [11] leads to the following payoffs for the home team:

$$\begin{aligned} h_{(0)}^{1,1} &= 1 - 2nr_H - nf + 3c(m-2) + dm(r_A - r_H) + s_H\{(m-1)(2-s_H) + nf - b(m-1) \\ &\quad + c(8-5m) - d(r_A - r_H)\} + 2s_Hs_A(2-m) + s_A\{2 + s_A - ne + b + c(8-3m) \\ &\quad - d(m-1)(r_A - r_H)\} \end{aligned}$$

$$\begin{aligned} h_{(0)}^{1,2} &= h_{(0)}^{1,1} + de(m - s_H + s_A - ms_A) + c\{m - 3 - 2ne + nf + s_H(5-3m) \\ &\quad + b(3-m) + c(24-11m) + s_A(4-m) - d(r_A - r_H + f)(2m-1)\} \end{aligned}$$

$$\begin{aligned} h_{(0)}^{1,3} &= h_{(0)}^{1,1} + df(m - s_H + s_A - ms_A) + b\{2m - 3 - ne + 2nf + s_H(4-3m) \\ &\quad + b(3-2m) + c(24-13m) + s_A(5-2m) - d(r_A - r_H + f)(m+1)\} \end{aligned}$$

$$h_{(0)}^{1,4} = h_{(0)}^{1,2} + h_{(0)}^{1,3} - h_{(0)}^{1,1}$$

$$\begin{aligned} h_{(0)}^{2,j} &= h_{(0)}^{1,j} + fz_H + c\{2m - ne + 2nf - 3 + s_H(4-3m) + b(3-2m) + c(24-13m) \\ &\quad + s_A(5-2m) + d(r_H - r_A + f)(1+m)\} \end{aligned}$$

$$\begin{aligned} h_{(0)}^{3,j} &= h_{(0)}^{1,j} + ez_H + b(m - 2ne + nf - 3 + s_H(5-3m) + b(3-m) + c(24-11m) \\ &\quad + s_A(4-m) + d(r_H - r_A + f)(2m-1)\} \end{aligned}$$

$$h_{(0)}^{4,j} = h_{(0)}^{2,j} + h_{(0)}^{3,j} - h_{(0)}^{1,j} \text{ for } j=1, \dots, 4$$

where z_H is defined as in [12]. For the away team, $a_{(0)}^{i,j}$ is obtained by replacing s_A, s_H, r_A and r_H with s_H, s_A, r_H and r_A , respectively, and z_A (similarly defined) with z_H , in the expression for $h_{(0)}^{j,i}$.

References

- Andreoni, J. and J. H. Miller, (1993), Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence, *Economic Journal*, 103, 570-585.
- Banerjee, A. N., Swinnen, J. F.N., and A. Weersink, (2007), Skating on Thin Ice: Rule Changes and Team Strategies in the NHL, *Canadian Journal of Economics*, 40, 493-514.
- Belot, M., Bhaskar, V., and J. van de Ven, (2010), Promises and Cooperation: Evidence from a TV Game Show, *Journal of Economic Behavior and Organization*, 73, 396-405.
- Bereby-Meyer, Y. and A. E. Roth, (2006), The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation, *American Economic Review*, 96, 1029-1042.
- Burton-Chellew, M. N., and S. A. West (2012), Correlates of Cooperation in a One-Shot High-Stakes Televised Prisoners' Dilemma, *PLoS ONE* 7, e33344. <https://doi.org/10.1371/journal.pone.0033344>.
- Bigoni, M., Casari, M., Skrzypacz, A., and G. Spagnolo, (2015), Time Horizon and Cooperation in Continuous Time, *Econometrica*, 83, 587-616.
- Buzzacchi, L. and S. Pedrini, (2014), Does Player Specialization Predict Player Actions? Evidence from Penalty Kicks at FIFA World Cup and UEFA Euro Cup, *Applied Economics*, 46, 1067-1080.
- Cabral, L. M. B., (2002), Increasing Dominance with No Efficiency Effect, *Journal of Economic Theory*, 102, 471-479.
- Chiappori, P. A., Levitt, S., and T. Groseclose, (2002), Testing Mixed-Strategy Equilibria when Players are Heterogeneous: The Case of Penalty Kicks in Soccer, *American Economic Review*, 92, 1138-1151.
- Dal Bó, P., (2005), Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games, *American Economic Review*, 95, 1591-1604.
- Dal Bó, P., and G. R. Fréchette, (2011), The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence, *American Economic Review*, 101, 411-429.
- Dal Bó, P., and G. R. Fréchette, (2015), Strategy Choice in the Infinitely Repeated Prisoners' Dilemma, <https://ssrn.com/abstract=2292390> or <http://dx.doi.org/10.2139/ssrn.2292390>.
- Duffy, J., and J. Ochs, (2009), Cooperative Behavior and the Frequency of Social Interaction, *Games and Economic Behavior*, 66, 785-812.

- Dawson, P., Dobson, S., Goddard, J., and J. O. S. Wilson, (2007), Are Football Referees Really Biased and Inconsistent? Evidence on the Incidence of Disciplinary Sanction in the English Premier League, *Journal of the Royal Statistical Society Series A*, 170, 231-250.
- Dixon, M. J., and M. E. Robinson, (1998), A Birth Process Model for Association Football Matches, *Journal of the Royal Statistical Society Series D*, 47, 523-538.
- Dobson, S., and J. Goddard, (2010), Optimizing Strategic Behaviour in a Dynamic Setting in Professional Team Sports, *European Journal of Operational Research*, 205, 661-669.
- Embrey, M., Fréchette, G. R., and S. Yuksel, (2017), Cooperation in the Finitely Repeated Prisoner's Dilemma, <https://ssrn.com/abstract=2743269>.
- Fréchette, G. R., and S. Yuksel, (2017), Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination, *Experimental Economics*, 20, 279-308.
- Friedman, D., and R. Oprea, (2010), A Continuous Dilemma, Mimeo, Economics Department, University of California, Santa Cruz.
- Garicano, L., and I. Palacios-Huerta, (2005), Sabotage in Tournaments: Making the Beautiful Game a Bit Less Beautiful, Center for Economic Policy Research Discussion Paper no. 5231.
- Hvide, H. K., (2002), Tournament Rewards and Risk Taking, *Journal of Labor Economics*, 20, 877-898.
- Hvide, H. K., and E. G. Kristiansen, (2003), Risk Taking in Selection Contests, *Games and Economic Behavior*, 42, 172-179.
- Kagel, J. H., and P. McGee, (2016), Team versus Individual Play in Finitely Repeated Prisoner Dilemma Games, *American Economic Journal: Microeconomics*, 8, 253-276.
- Levitt, S., (2004), Testing Theories of Discrimination: Evidence from Weakest Link, *Journal of Law and Economics*, 47, 431-452.
- Levitt, S., and J. A. List, (2007), What do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?, *Journal of Economic Perspectives*, 21, 153-174.
- List, J. A., (2006), Friend or Foe? A Natural Experiment of the Prisoner's Dilemma, *Review of Economics and Statistics*, 88, 463-471.
- Lugovsky, V., Puzzello, D., Sorensen, A., Walker, J., and A. Williams, (2017), An Experimental Study of Finitely and Infinitely Repeated Linear Public Goods Games, *Games and Economic Behavior*, 102, 286-302.

- Mao, A., Dworkin, L., Suri, S., and D. J. Watts, (2017), Resilient Cooperators Stabilize Long-Run Cooperation in the Finitely Repeated Prisoner's Dilemma, *Nature Communications*, 8, article number:13800.
- Mengel, F., (2014), Learning by (Limited) Forward Looking Players, *Journal of Economic Behavior and Organization*, 108, 59-77.
- Oberholzer-Gee, F., Waldfogel, J., and M. W. White, (2010), Friend or Foe? Cooperation and Learning in High-Stakes Games, *Review of Economics and Statistics*, 92, 179-187.
- Palacios-Huerto, I., (2003), Professionals Play Minimax, *Review of Economic Studies*, 70, 395-415.
- Palomino, F., Rigotti, L., and A., Rustichini, (2000), Skill, Strategy, and Passion: An Empirical Analysis of Soccer, Tilberg University, Discussion Paper, no. 129.
- Ridder, G., Cramer, J. S., and P. Hopstaken, (1994), Down to Ten: Estimating the Effect of a Red Card in Soccer, *Journal of the American Statistical Association*, 89, 1124-1127.
- Van den Assem, M. J., Van Dolder, D., and R. H. Thaler, (2011), Split or Steal? Cooperative Behavior when the Stakes are Large, Mimeo, Erasmus University of Rotterdam.
- Walker, M., and J. Wooders, (2001), Minimax Play at Wimbledon, *American Economic Review*, 91, 1521-15