

Fischer, Kai; Schmal, W. Benedikt

Article — Published Version

Pricing in response to new information: The case of betting markets

Economic Inquiry

Suggested Citation: Fischer, Kai; Schmal, W. Benedikt (2024) : Pricing in response to new information: The case of betting markets, Economic Inquiry, ISSN 1465-7295, Vol. 63, Iss. 1, pp. 236-264, <https://doi.org/10.1111/ecin.13258>

This Version is available at:

<https://hdl.handle.net/10419/313802>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

ORIGINAL ARTICLE

Pricing in response to new information: The case of betting markets

Kai Fischer¹ | W. Benedikt Schmal^{2,3} 

¹Düsseldorf Institute for Competition Economics (DICE) at Heinrich Heine University Düsseldorf, Düsseldorf, Germany

²Economic Theory Group, Ilmenau University of Technology, Ilmenau, Germany

³Department for Management, Strategy, and Innovation (MSI), KU Leuven, Leuven, Belgium

Correspondence

W. Benedikt Schmal, Economic Theory Group, Ilmenau Ehrenbergstr. 29, Ilmenau 98693, Germany; Department for Management, Strategy, and Innovation (MSI), KU Leuven, Namsestraat 69, Leuven 3000, Belgium.
Email: wolfgang-benedikt.schmal@tu-ilmenau.de

Funding information

Deutsche Forschungsgemeinschaft, Grant/Award Number: 235577387/GRK1974

Abstract

Markets are information aggregators. But how do they incorporate new data into their pricing? We examine the response of prediction markets to a novel information shock in a quasi-natural experiment: How did the absence announcements of elite soccer players influence the betting odds of affected matches? Analyzing the first four statistical moments of 117,174 odds from 32 bookmakers, we identify initial inertia followed by a lagged reaction that we cannot reason with learning. Our findings raise questions about how bettors and bookmakers incorporate new information into their beliefs. It has broader implications regarding information processing in markets.

KEYWORDS

belief updating, betting markets, betting odds, forecasting errors, information shocks, prediction markets, pricing

JEL CLASSIFICATION

D83, D84, G14, D22, D53

1 | INTRODUCTION

The superiority of pooled over individual information processing has been discussed in various settings.¹ It originates from the fundamental understanding of markets as large, decentralized aggregators of information as outlined by Hayek (1945) and formalized for expectations by Muth (1961). Financial markets are compelling in this context as they get close to a frictionless market environment with perfect competition and optimal information aggregation

Abbreviations: COVID, coronavirus disease; CVR, coefficient of variation; DiD, difference in differences; EGBA, European Gaming and Betting Association; Eq., equation; EU, European Union; EUR, Euro (currency); Fig., figure; Hyp., hypothesis; Inf., infection; LHS, left hand side; OLS, ordinary least squares; p., page; PDF, probability density function; RHS, right hand side; SUTVA, stable unit treatment value assumption; UK, United Kingdom; USD, US Dollar.

Managing Editor: Rob Simmons

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). Economic Inquiry published by Wiley Periodicals LLC on behalf of Western Economic Association International.

(Grossman, 1976, 1978; Hellwig, 1980). In such an environment, we study how prices of established future contracts with a finite time horizon incorporate new information. We rely on the high-stakes prediction market of sports betting to benefit from the existence of final outcomes for betting on finite events (i.e., matches), which we can use to evaluate prior predictions.

We exploit the global COVID-19 pandemic as a natural experiment. It allows us to unambiguously identify the absence of athletes due to coronavirus infections. These are exogenous information shocks to both bookmakers and bettors. We study professional men's soccer in Europe, one of the world's most sought-after entertainment markets. The elite leagues implemented a rigorous testing scheme for the virus. Infections had to be made public immediately—unlike most sicknesses or absences. It allows for precise separation of “treated” and “control” matches. Infected players were obliged to quarantine for more than a week following a positive test. They were missing at least one match, which constitutes our information shock.

Our setting applies to many more markets aside from betting than analyses of shocks in, for example, live betting environments that occur over an actual match or race and usually study only one betting firm (Choi & Hui, 2014; Crosson & Reade, 2014). By investigating arguably all firms and, hence, a whole market, our results are fairly general and may be reflected in many instances of incorporating new information into pricing and purchase decisions.

The analysis builds upon odds from 32 bookmakers covering most of the sports betting market in Europe. The market for sports betting generated EUR 14.6 billion in revenues in 2021, which accounts for 40% of the total online gambling market.² The natural quasi-experiment in this large prediction market is a high-stakes environment, as bookmakers' profits crucially depend on the correct forecasting of match outcomes.

We study 1372 matches in the premier divisions of men's soccer in Germany and Italy during the seasons 2019/2020 and 2020/2021, covering a period from autumn 2019 to summer 2021. Within this comparatively short time window, COVID-19 equipped us with 257 positive cases, of which we can unambiguously assign 233 cases to the infected players. With an accuracy of 90.7%, this meticulously gathered player-level infection data enable a clear distinction between treatment and control groups and hardly suffer from any dark figures.

The high number of bookmakers and odds enables us to analyze the heterogeneity and dispersion in an entire market instead of studying a single bookmaker's behavior.³ It allows us to conduct a manifold analysis of market-level uncertainty and aggregated information incorporation in response to a reasonably general shock, as player absenteeism is an everyday business in professional team sports but has yet to be seamlessly explored. We focus on the first four moments of the odds distribution of each matchday. Our differences-in-differences framework utilizes the timing and announcement of infections as causal identification. We propose four behavioral hypotheses and test them empirically. The ex-post benchmark would be a null effect, as we cannot find any effect of a player's absence on match outcomes. Nevertheless, there seem to exist distorted beliefs, especially in the later stages of the pandemic: While we do not identify any effect on the odds distribution among early infections, later infections show a significant increase in the affected team's average odds and their dispersion and a jump in the absolute skewness of odds. It is a strong hint that either bettors, bookmakers, or both react differently to our data-driven (ex-post) benchmark.

According to Snowberg et al. (2013), prediction markets work best if their mechanism successfully aggregates information, truthful disclosure of the information is rewarded, and long-term incentives exist to invest in improved information gathering. For sports betting, these criteria are mostly satisfied. Furthermore, these markets tend to process new information quickly. This ties in with the “Hayek hypothesis” (Hayek, 1945; Smith, 1982). It states that market participants—only equipped with limited knowledge of the overall environment—can create an efficient outcome using the market process alone.⁴ The hypothesis holds under five conditions: The number of traders is small, and they can set prices. They only have incomplete information on the market environment and can or cannot have rational expectations of the market environment. Last, no central authority coordinates the market (Al-Ubaydli et al., 2022). We consider these conditions mostly satisfied here, even though a small number of ‘traders’ only applies to bookmakers while plenty of bettors are in the market.

Moreover, our work is essentially a test of the semi-strong efficient market hypothesis (Fama, 1970). It states that prices in financial markets reflect all publicly available information. As player absenteeism is known to the public and the match outcomes of past matches, one would expect market participants to incorporate new absence announcements, assuming semi-strong market efficiency correctly.

The subsequently identified inertia questions how well prediction markets incorporate new information. Berg et al. (2008) show for presidential election forecasting in the US that forecasts early in the election campaign outperform common polls, implying that prediction markets are compelling for long-run forecasts. Deck and Porter (2013) survey

lab experiments on the strength of markets in information aggregation and find evidence for the baseline hypothesis suggested by Hayek.

Sports betting is widely used for a better understanding of biases and anomalies in asset pricing and financial markets more generally (see, e.g., Avery & Chevalier, 1999; Brown & Yang, 2016; Feddersen et al., 2017; Moskowitz, 2021; Quandt, 1986; Ramirez et al., 2023; Thaler & Ziemba, 1988; Wolfers & Zitzewitz, 2004). An important example for research on behavioral anomalies is the “longshot bias:” Bettors overvalue the chances of winning for an unlikely event, usually the win of an “underdog” (see, e.g., Meyer & Hundtofte, 2023; Whelan & Hegarty, 2023; Whelan, 2024). Choi and Hui (2014) find for sports betting that unsurprising events lead to an overly passive reaction not fully reflecting the extent of the news. In contrast, there is overreaction after the arrival of what they call surprising events. Their sole incident of interest, however, is a goal scored. Angelini et al. (2022), Croxson and Reade (2014), and Ötting et al. (2023) investigate the effect of goals on odds in live betting and find biased reactions to another goal, which constitutes new information regarding the expected final match outcome. As live betting appears limited to a narrow set of potential applications elsewhere, we study bookmakers’ odds setting along the whole distribution of odds. In contrast, previous studies focused on one firm.

Winkelmann et al. (2024) look at betting odds in the long run and find that mispricing disappears over time and markets become efficient. It does not contradict our work as we investigate the short run after the advent of new information. Brown et al. (2019) find the release of a new US presidential election poll to trigger more activity on a betting platform and more noise in the valuations. We provide novel evidence of inertia in response to hardly quantifiable information. Williams and Reade (2016) relate political betting with social media activity and find a lagged response to a shock. It corresponds to our findings regarding late infection announcements. Brown (2014) identifies limitations in the information processing ability as one driver of mispricing in betting markets. It should not so much apply to bookmakers who are used to analyze new information but may affect bettors in their belief formation.

The remainder of this paper is structured as follows. Section 2 explains the data we use and the sports betting market. Section 3 describes our empirical strategy. Section 4 states the core behavioral hypotheses we investigate and presents the benchmark case. We present our findings in Section 5. Section 6 concludes.

2 | DATA AND BACKGROUND ON SPORTS BETTING

The closing odds of 32 bookmakers, obtained from *oddsportal*, build the foundation of our analysis.⁵ The online gambling and sports betting market is highly entangled because many betting brands only operate in selected markets or are sub-licensed to other companies for legal reasons. Members of the main European business association EGBA are *bet365*, *Betsson*, *Entain*, *Flutter*, *kindred*, and *William Hill*. All firms usually run several gambling and betting brands. We have at least one betting platform with odds for each firm in our sample. Our sample also contains odds from *Leonbet*, particularly engaged in Russian-speaking countries, and *Dafabet*, focused on Asian markets.

We cover the seasons 2019/2020 and 2020/2021 of the elite leagues in Germany (Bundesliga) and Italy (Serie A).⁶ Such a season usually lasts from late summer until spring of the next year. Hence, we have COVID-19 infections in the later stages of the 19/20 season and the whole 20/21 season, with a particular peak in winter 20/21. The Bundesliga consists of 18 teams, and the Serie A of 20 teams. It leads to a total number of 1372 matches. To obtain a balanced team-matchday panel, we double the data set as the 1X2 odds are defined by *match* and not by *team*.⁷ We cluster the standard errors at the match level to account for it. In total, we use 117,174 single odds.⁸

Technically, we use “1X2 odds,” that is, odds for a win of the home team (1), a draw between both teams (*X*), or a win of the away team (2). These odds are in decimal-style; for any odds ξ holds $\xi \in (1, \infty)$. For any wager ω , the payout is $\xi\omega$ and profit is $\pi = \xi\omega - \omega$ in case of a successful bet and $\pi = -\omega$ otherwise. Higher odds lead to a higher payoff in case of success but are related to a lower probability of success. We code “1” and “2” odds as treated in the case that an infection happens. We separate the infection data at the median player-level infection date to distinguish between early and later infections. We use the closing odds of each bookmaker, that is, the last odds published before a match is about to begin.

Our analysis relies on the accurate identification of treatment and control observations. Here comes COVID-19 into play: Only for this virus disease a rigorous testing pattern and public announcements of new infections existed. Other common sicknesses, such as chills or gastrointestinal diseases, are often announced on short notice before a match or not at all. Hence, one cannot detect an information shock for them. COVID-19 provides a sublime exception, so we exploit it as an information shock. We use a detailed list of infections meticulously gathered by Fischer et al. (2022).

They identified 233 players that occurred ahead of 131 matches.⁹ Matching players to their teams enables us to distinguish between affected and unaffected teams.

Many players got infected simultaneously, especially in the later stages of the pandemic. We assume a constant marginal effect of a new infection *event*—so, the bundle of at least one infection before a team's next match—on betting odds. While it may be a significant difference whether, say, one or five team players are newly missing, it is reasonable to assume that the bookmakers should have picked up the absence of players already missing for some time. For example, betting odds for the upcoming match should capture the absence of a previously infected player who has been missing for 3 weeks. Any changes caused by an *additional* infection (and, thus, absence) should be solely related to the new announcement. Hence, only the *additional* absence should affect the betting odds. The whole dataset is publicly available on openIPSCR (Fischer & Schmal, 2024).

3 | EMPIRICAL SETUP

Regression Design: We analyze the announcement of a missing player on the betting odds of the respective teams and their opponents. This event has two effects: First, the choice set of the team manager decreases as there are fewer players to choose from for the squad on the pitch, which weakly lowers quality. Second, due to an absence, the squad on the pitch might need an adjustment if the absentee is a regular player. Our reduced-form estimation for betting odds aggregates both issues and captures the resulting uncertainty of how a team's ability might change because of this. Empirically, we apply a difference-in-difference design that uses the infection date as information shock that should affect the belief formation of bookmakers and bettors. We define

$$f(\xi_i)_{tm} = (T_{tm} \times S_m)' \beta_T + Z' \zeta + \epsilon_{tm}. \quad (1)$$

The dependent variable is the specific moment measure of the closing odds ξ of i bookmakers for team t in match m —either the mean, the standard deviation (both logarithmic), the skewness, or kurtosis (both in levels, by construction) across all bookmakers i in the market.¹⁰ β_T is the coefficient vector for the treatment effect. S_m is a binary indicator, whether it is an early or a late infection. T_{tm} is the treatment indicator (a team's first match after at least one infection is reported). We separate these two groups at the median infection date. Intuitively, a distinction between seasons might suggest itself. However, as Figure A3 in the appendix shows, the number of infections in season 2019/2020 is comparatively low. Therefore, we use the *median* as the main measure of distinction, but our results are robust to variations.¹¹ Z is a vector of fixed effects, namely team \times season, opponent team \times season and matchday \times season. The vector ζ , accordingly, describes the estimated coefficients. ϵ_{tm} is the idiosyncratic error term. To extend our analysis, we also apply a dynamic difference-in-differences design in the form of an event study, which we model as

$$f(\xi_i)_{tm} = \sum_{\tau=\underline{k}, \tau \neq -1}^{\bar{k}} \beta_\tau T_{tm,\tau} + Z' \zeta + \epsilon_{tm}. \quad (2)$$

Outcome Variables: The first and last bins \underline{k} and \bar{k} are binned endpoints, encompassing all other observations outside the effect window. All other terms are equivalent to those in Equation (1). The first two moments of the distribution of odds by team and matchday are

$$\mu_m = \frac{1}{N_m} \sum_{i=1}^{N_m} (\xi_{im}) \quad (3)$$

for the mean and

$$\sigma_m = \sqrt{\frac{1}{N_m} \sum_{i=1}^{N_m} (\xi_{im} - \mu_m)^2} \quad (4)$$

the standard deviation as the square root of the variance for the second moment.¹² We use the Fisher-Pearson coefficient of skewness to work with the third moment of a distribution. It describes the third moment around the mean and

can be considered the oldest and most established measure of skewness (Doane & Seward, 2011; Singh et al., 2019). Applied to the odds ξ_i per matchday m , this becomes

$$\gamma_m = \frac{1}{N_m} \sum_{i=1}^{N_m} \left(\frac{\xi_{im} - \bar{\xi}_m}{\hat{\sigma}_m} \right)^3. \quad (5)$$

For all moments, N_m captures the number of bookmakers providing odds at matchday m . While the importance of skewness for individuals in financial markets is widely recognized—see, for example, Alderfer and Bierman (1970), Golec and Tamarkin (1998), and Hirshleifer (1965); more recently Dertwinkel-Kalt and Köster (2019) and Ebert (2015)—the third moment has been utilized to a lesser extent as an outcome of markets. As pointed out by Mao (1970), there are notable differences in the perception of two gambles with the same mean and variance but different skewness values. Our paper attempts to exploit this on a market level instead of individuals.

Lastly, we look at the fourth moment of the betting odds distributions, namely the kurtosis. Analogously to the formula of skewness presented in Equation (5), the kurtosis κ is defined in our setting as

$$\kappa_m = \frac{1}{N_m} \sum_{i=1}^{N_m} \left(\frac{\xi_{im} - \bar{\xi}_m}{\hat{\sigma}_m} \right)^4. \quad (6)$$

The kurtosis is usually interpreted relative to the normal distribution, which has $\kappa = 3$, distributions with $\kappa < 3$ are considered platykurtic and have less extreme and fewer outliers than the normal distribution. In contrast, distributions with $\kappa > 3$ are called leptokurtic and have more extreme outliers and a higher number of them (see, e.g., Balanda & MacGillivray, 1988).

Identification: To disentangle the effect of the public announcement of new infections, we compare the betting odds of infected players' teams before and after their positive test results with the evolution of outcomes of non-infected players and their respective teams. We apply a two-way fixed-effects difference-in-differences estimation that controls for variation over time and across clubs.¹³

Several assumptions need to hold for valid results in this empirical setting. Within our difference-in-differences design without dynamic effects, we need parallel trends for the treatment and control group without COVID-19. There is no conceivable cause for the diverging evolution of betting odds without COVID-19. For the dynamic event study setting in Equation (2), it corresponds to the requirement that treatment cannot predict outcomes prior to treatment. Our event study plots show flat pre-trends, indicating that the assumption of parallel trends is satisfied.

Furthermore, we assume that there exists no self-selection of players in the treatment. In contrast to other injuries, for which clubs can strategically decide on what they want to announce at which time, the mandatory and transparent COVID-19 testing schemes diminish this concern. Also, we consider infections as not anticipated in the short-run—neither by players, teams, or bookmakers. As Table A2 in the Appendix illustrates, we can proceed on the assumption that the assignment of new infections and, by that, absence announcements are random. No statistical relation exists to the league, earlier announcements, previous performance measured in points gained, or lagged moments of the odds.

Eventually, a difference-in-differences estimator requires that the treatment only causes partial equilibrium effects, or else we need the stable unit treatment value assumption to hold. While Fischer et al. (2022) find spillover effects in players' performance on the pitch, there is no reason why betting odds of unaffected matches should change. It would be theoretically conceivable that a COVID-19 infection raises the market value of uninfected players in the same position (e.g., striker or midfielder) because the supply of healthy players in that particular position decreases. However, European soccer is so large that the comparatively low number of infections is unsuspicious of causing general equilibrium effects besides some superstar players, which, on the other hand, are hardly statistically significant. Also, almost no player trading happens throughout the run of a season. Thus, we consider the assumptions required for causal estimates within the difference-in-difference design as satisfied.

To ensure that our regression setup identifies the effect of an actual infection, we need to ensure that COVID-19 infections do not spill over within teams to a large extent. Figure 1 shows the effect of one infection in a team on the number of reported infections in the team in the subsequent days. While there is evidence for positive spillovers in the very short run, the effects are limited in size, with less than 0.1 additional infections in the first week after a reported




FIGURE 1 Effect of an infection on new infections within a team. This figure plots the OLS coefficients of a regression of team-level infections on lags of an infection dummy. Standard errors: Heteroskedasticity-robust and clustered on the match level. 90% (blue lines) and 95% confidence bands (black extensions) are shown.

infection. The effect is also limited to the first few days after infection, so the uncertainty about new infections mostly disappears until the first post-infection match. We take this as indicative evidence that our analysis later on likely does not capture the pure uncertainty of potentially many upcoming new infections.

4 | BEHAVIORAL HYPOTHESES

To formalize the behavioral responses of betting markets to the information “shock” of a new absence, one should first consider the basic functioning of betting markets. Figure 2 sketches this: It is reasonable to assume that betting companies solely rely on information to form their betting odds and not on the intuitions of their employees. In contrast, sports betting is closely related to gambling, so it can be considered an entertaining and emotional task for many bettors. They are likely prone to incorporate “gut feelings” such as emotional attachment to a club, behavioral biases such as the hot hand fallacy, or simply subjective evaluations of teams and players into their formation of expectations and computation of winning probabilities. However, information also affects intuition in the sense that it might shape biases. Crucial for understanding betting markets is that bettor behavior also drives bookmaker behavior, consequentially also betting odds and, eventually, the intuition of the group of buyers (Levitt, 2004).

The rationale behind this is that betting companies need to adjust their portfolio of closed bets since bets imply a payment obligation in case the buyer is right with their prediction. If non-outcome-related drivers affect the composition of bets, betting companies need to adjust for their expected outcomes. One way of doing that is by changing the odds to balance their portfolio to some extent. It implies that the odds adjustment could be driven not only by the betting firms’ behavioral biases but also by the bettors’ anticipated behavioral biases that lead to an adjustment of the odds after a recently announced player’s absence. Regarding new cases, the only reasonable behavior of bettors besides remaining indifferent would be shying away from bets on a win for the affected team. In turn, the betting companies would react by adjusting the odds for the affected team upwards and the opponent team downwards.

However, this effect should occur predominantly in the earlier stages of the pandemic, when the emotional response and media attention were exceptionally high. In addition to traditional bookmakers, we can exploit data from “Betfair exchange,” which is not a standard bookmaker but—as the name suggests—an exchange. There is no central company offering bets to a crowd of customers, but it is a genuine market with many sellers and buyers of bets. Thus, these users do not face the same challenge of large-scale portfolio balancing. Figure A4 in the Appendix shows the average odds of




FIGURE 2 Stylized relations between bettors, betting companies, and information.

this exchange to react in the same way as the bookmakers' odds. Hence, we do not suspect the steering of (naive) consumers via portfolio balancing to be a relevant driver in this domain, as the betting market as a whole appears to be distorted. We also compute the bookmakers' profit margins as defined by Shin (1993) to check for changes in the market characteristics. Figure A5 in the Appendix highlights the absence of margin changes due to the announcement of new absences. We also rule out changes in the odds due to variations in competition intensity.

Players missing matches is a widespread phenomenon for clubs and the public. Athletes get injured, get ordinary infections such as a cold or influenza, or do not play for disciplinary or private reasons. Henceforth, bookmakers should be aware of that problem, and one would expect that they have fitted their own predictive models accordingly. Based on these considerations, we propose four different hypotheses on the markets' reaction to the advent of new information on player absences using the four moments of the underlying theoretical distributions of betting odds. There are three potential effects for each moment: a significant increase, decrease, or no change. Together with four moments considered in total, this results in 4^3 potential combinations. We narrow our focus to the following four hypotheses that we consider the most reasonable:

Hypothesis 1. Null Effect

$$\Delta E_m[\xi_i] = 0 \wedge \Delta V_m[\xi_i] = 0 \wedge \Delta \gamma_m[\xi_i] = 0 \wedge \Delta \kappa_m[\xi_i] = 0$$

Hypothesis 1 states that there is no observable effect in the betting odds distribution at all. It could happen if the new information appears worthless to the agents and no reaction is considered necessary. In that case, none of the four moments should react.

Hypothesis 2. Educated Adjustment

$$\begin{cases} \Delta E_m[\xi_i] \neq 0 \wedge \Delta V_m[\xi_i] = 0 \wedge \Delta \gamma_m[\xi_i] = 0 \wedge \Delta \kappa_m[\xi_i] = 0 & \text{valuable information} \\ \text{Hypothesis 1} & \text{valueless information} \end{cases}$$

Second, there could be an educated adjustment as described by Hypothesis 2. The agents recognize new information and its meaning and incorporate it properly into their expectations. If there is no change in uncertainty regarding the size and/or direction of the effect caused, there should be only a change in the mean to reflect the changed state of the world but no change in variance, skewness, or kurtosis as there is no change in uncertainty. This hypothesis collapses to the null effect hypothesis (Hyp. 1) if infections do not alter match results.

Hypothesis 3. Uneducated Uncertainty

$$\Delta E_m[\xi_i] = 0 \wedge \Delta V_m[\xi_i] \geq 0 \wedge \Delta \kappa_m[\xi_i] \geq 0$$

Third, subsequent to the arrival of new information, agents might be unsure how to incorporate it into their prediction but expect the information to be meaningful. Hypothesis 3 captures this. As in Hyp. 1, there is no change in the mean. It is now reasoned by the uncertainty about *how* the event affects the predicted outcomes. Contrary to Hyp. 1, the second moment increases as a jump in uncertainty urges the agents to adjust their predictions. The same holds for the kurtosis. Regardless of its initial value, a weakly positive change implies the distribution to become less platykurtic or else more leptokurtic. Hence, more and more extreme outliers should occur. It is unclear whether the skewness of

the predictions should change as increased uncertainty might reduce the skewness of a previously skewed distribution or else increase the skewness of an unskewed distribution. Therefore, this hypothesis remains agnostic regarding potential changes in the third moment.

Hypothesis 4. Rally round the flag

$$\Delta E_m[\xi_i] = 0 \wedge \Delta V_m[\xi_i] \leq 0 \wedge \Delta \gamma_m[\xi_i] \neq 0 \wedge \Delta \kappa_m[\xi_i] \leq 0,$$

$$\text{where } \Delta \gamma_m[\xi_i] \begin{cases} > 0 & \forall \gamma > 0 \\ < 0 & \forall \gamma < 0 \end{cases}$$

Hypothesis 4 proposes a potential “rally round the flag” effect under uncertainty. It is well known in conflict situations, in which citizens tend to gather around their leaders much more strongly than they might do in more peaceful times (Waltz, 1967). As there is no natural “leader” in the betting market, we consider the skewness parameter instead as a measure of concentration beyond the standard deviation that captures dispersion. However, this only holds conditional on the mean and the standard deviation to remain stable. It is equivalent to Hyp. 3 in so far that the mean is not supposed to change because there is no clear information on how the infections impact the respective outcomes (match results and individual quality).

Contrary to Hyp. 3, the measure of dispersion either decreases or else remains unaffected. It should also be visible in the kurtosis, which should become less leptokurtic and more platykurtic, implying fewer outliers. However, a change in skewness implies a movement toward a higher concentration conditional on holding mean and variance fixed. As a change in γ might also reduce skewness, we specify that $\Delta \gamma_m[\xi_i]$ is always related to increasing skewness for Hyp. 4 to hold. The last hypothesis is necessary because even without a change in dispersion, there might be a change in concentration, that is, the first two moments do not change, but the third.

In general, more combinations of changes in the four moments are conceivable. However, we consider these four scenarios as the most reasonable behavioral reactions. As we estimate a reduced-form model, we cannot reveal evidence for a specific theory, but we aim to get a deeper understanding of the existing behavioral patterns. We try to understand the path toward a new equilibrium by looking at the development of adjustment over time. We look at this process from the advent of novel information and study how betting odds react.

Benchmark: To understand (distorted) belief formation, we first need to evaluate what would have been ex-post rational to expect for the market participants. Figure 3 presents the event study plots for two central match

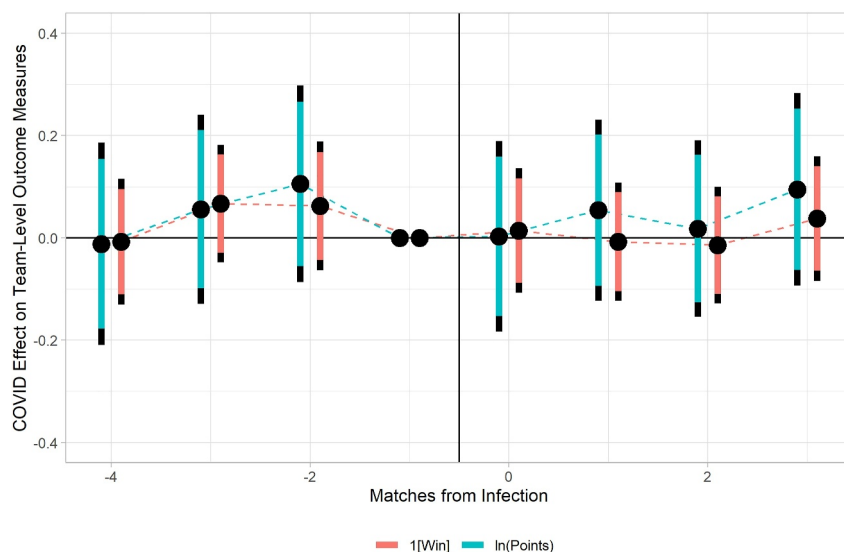


FIGURE 3 Dynamic effect on points and goals scored for a team with a new absence. This figure plots the dynamic OLS coefficients β_T for the probability of winning (blue) and logarithmic points gained (red) by teams with new infections. Standard errors: Heteroskedasticity-robust and clustered at the match level. 90% (colored lines) and 95% confidence bands (black extensions) are shown. We use the inverse hyperbolic sine transformation for both outcomes to account for zero values.

outcomes in soccer as a benchmark. The estimates with blue confidence bands depict the probability of winning. Those with red bands show the points gained before and after a newly announced infection (and the absence of the affected player). While the flat pre-trends hint at treatment and control clubs being on the same trends, entirely insignificant estimates after a new infection are evidence of a lack of an absence effect on the team level. As the absence of players should be familiar to bookmakers, the knowledge on how to incorporate them into odds should be dispersed across agents in this market. Put differently, the odds should not react. However, this is *ex-post* reasoning. *Ex-ante* belief formation could differ without being irrational.

To ensure that the insignificant results are not driven by a lack of statistical power due to the decomposition of the effect into several bins, we also estimate the baseline DiD coefficients for the first post-infection match and a separation of that coefficient by early and late infections. This is shown in Figure 4.

The left panel displays the effect on the probability of winning a match as the dependent variable. The right panel depicts the effect on the number of points a team gains during a match. In both cases, the information shock of a newly missing player does not affect any of the measures. Furthermore, there is no difference between early and late infections. Hence, the only reasonable *ex-post* belief would be a null effect. There could be confusion at the beginning and some learning over time. So, some moments may be affected by early infections but certainly not by later ones.

5 | RESULTS

While the COVID-19 pandemic has caused global economic shock waves, sports betting remained surprisingly inert. There already exists literature that finds bookmakers insufficiently adjusted their odds for the effect of spectator lockouts on home advantage (Fischer & Haucap, 2022; Winkelmann et al., 2021). However, we take a more general approach as we study the effect of new absence announcements on the distribution of betting odds. It is essential to notice that we do *not* compute any COVID-19 effect but the overall effect of new information on odds.

5.1 | Changes in the four moments

Effect on average odds: We start with the first moment, that is, the average odds for a win in the match following the announcement of a new (additional) infection. The results are displayed in Figure 5. The effect crucially depends on

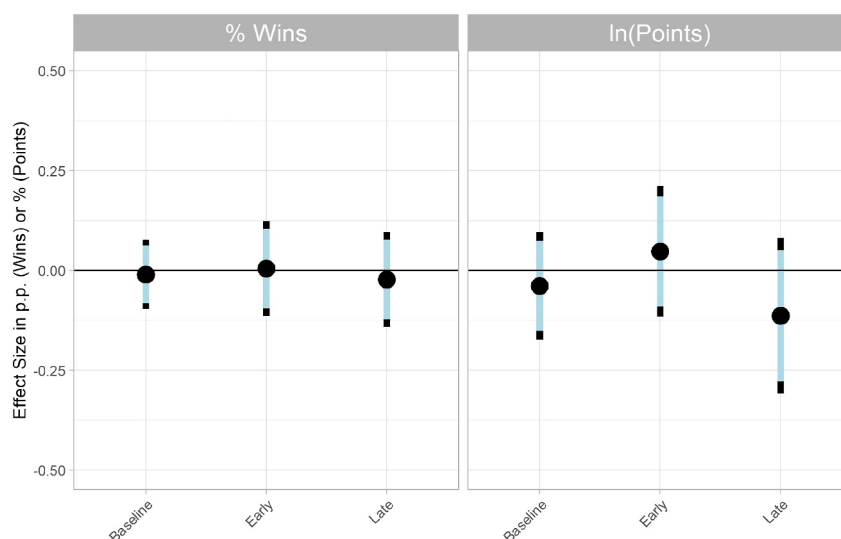


FIGURE 4 Match outcomes subsequent to a newly absent player. This figure plots the OLS coefficients β_T of the DiD regression following Equation (2) for the mean of the probability to win (left panel) and for the points gained (right panel). Standard errors: Heteroskedasticity-robust and clustered at the match level. 90% (blue lines) and 95% confidence bands (black extensions) are shown. For both outcomes, we use the inverse hyperbolic sine transformation to account for zero values but still be able to conduct a logarithmic transformation.

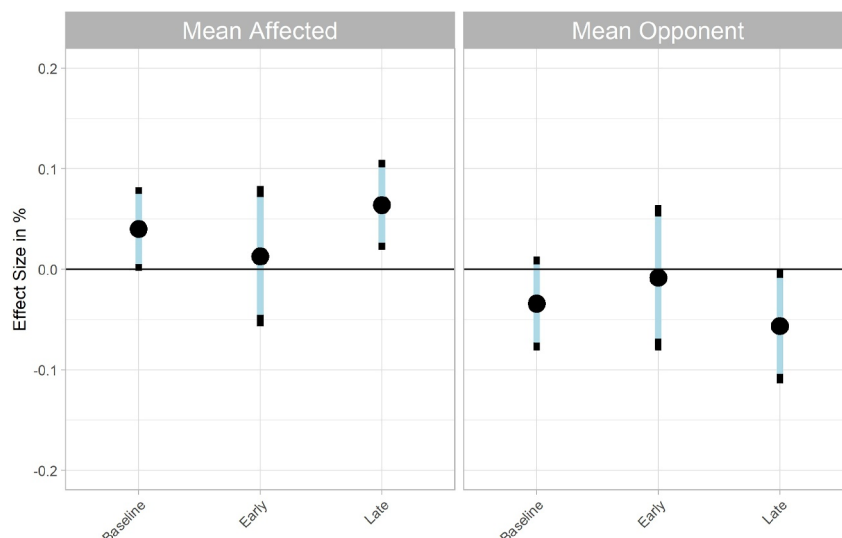


FIGURE 5 Effect on the average odds for a win after the announcement of a new absence. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the mean of the matchday odds for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions). Dependent Variable: Logarithmic Mean.

the timing of an infection. The left panel represents the coefficients for the affected team. The right one corresponds to the coefficients of the respective opponent team. Second, we distinguish between the baseline static DiD effect and a decomposition of the effect in early and late infections. We use the median infection as the cut-off.

Studying the average odds for a win of the affected team, we find no effect for early infections but a sharp positive effect for the second half of infections (significant even on the 1% significance level). However, they seem to remain surprisingly inert to this new information, especially because early infections created much more uncertainty and turmoil, which might also affect team performance. As higher odds imply a lower expected probability of the related match outcome, betting markets seem to be either inert or expect no effect of a COVID-19 infection on team performance in the first half of the affected matches but adjust their expectations toward a negative effect in the second half. Decomposing the effect on the mean odds for late infections further, we find that the effect occurs on the first match after the new infection (as shown in Figure 6). The flat pre-trends are indicative of the DiD setting's validity. For robustness, we further test the efficiency of the betting market using the approach of Angelini and De Angelis (2019). Hence, we regress the prediction error on the treatment.¹⁴ We present the results in Table A3 in the Appendix. We do not find significant mispricing following this approach. This is insightful for the potential learning process of the betting market over time. However, this approach does not allow us to study the second to fourth moment of the distribution.

Effect on the standard deviation: Figure 7 shows the main findings of the regressions regarding the standard deviation.¹⁵ Again, we distinguish between the baseline treatment effect and an additional timing decomposition. As discussed earlier, previous research has shown that new information leads to more activity. Hence, we expect a higher standard deviation after a new infection. This is the case—but only for the second half of the sample. Here, the standard deviation of set odds increases, a sign that the betting market does not precisely know how to incorporate the new information, but they make an attempt.¹⁶

However, this raises doubts about the hypothesis of learning. Suppose market participants learned something specific about the dynamics of COVID-19 infections on match outcomes. In that case, the variation in the odds setting should not increase except for the possibility that different companies have drawn different conclusions. However, Figures 3 and 4 in the “benchmark” section show hardly any outcome-related effect of newly missing players on the team level. Hence, betting odds on match outcomes should vary much less than we observe in the data. The dynamic event study in Figure 8 shows the equivalent to the dynamic effect for the mean odd: The first matchday after an infection announcement is significantly affected.

Effect on the skewness of the odds: The third pillar of our analysis is the skewness of the odds as defined in Equation (5) beforehand. As Table 1 shows, the overall skewness is centered approximately around zero but with a




FIGURE 6 Dynamic effect on the mean winning odds for late infections. This figure plots the dynamic OLS coefficients β_T of the DiD regression following Equation (2) for the log mean of the matchday odds for teams with new infections within the subsample of late infections. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

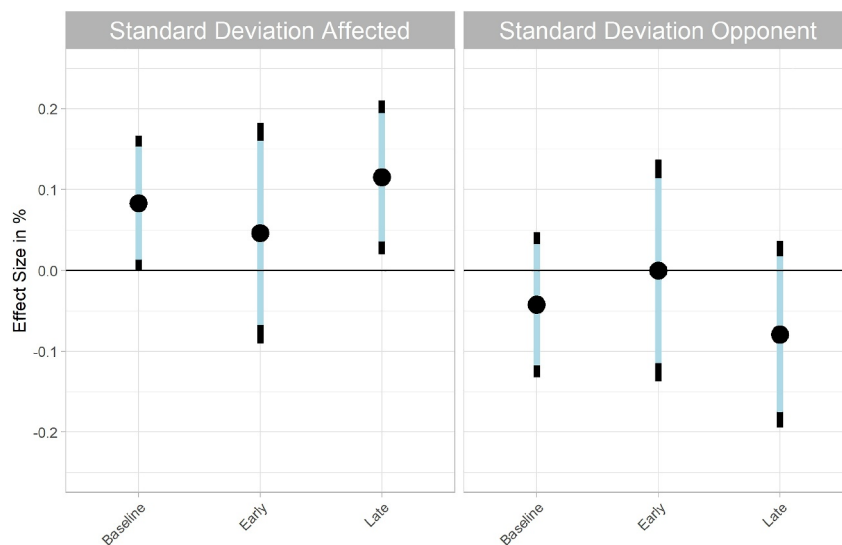


FIGURE 7 Effect on the odds' standard deviation after a new absence announcement. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the log standard deviation of the matchday odds for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

negative mean. It implies slight right-skewness of the betting odds. Distinguishing between early and late infections, one clearly sees that the odds of affected matches in the second half of the sample are much more left-skewed than in the first half ($E[\gamma_{late}] = -0.426$ and $E[\gamma_{early}] = -0.0626$).

We use the third moment of the odds distributions to investigate the odds movement and test for the already mentioned “rally round the flag” effect. As there is no natural “leader” in the betting market, we consider the skewness parameter as a measure of concentration beyond the standard deviation that captures dispersion. A change in skewness




FIGURE 8 Dynamic effect on the standard deviation for late infections. This figure plots the dynamic OLS coefficients β_T of the DiD regression following Equation (2) for the log standard deviation of the matchday odds for teams with new infections within the subsample of late infections. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

TABLE 1 Summary statistics on the skewness of betting odds.

Sample	Minimum	1th quartile	Median	Mean	3th quartile	Maximum
Early inf.	-1.0419	-0.5082	-0.1161	-0.0626	0.2990	1.4679
Late inf.	-2.4302	-0.7844	-0.5025	-0.4260	-0.0025	0.4774
No inf.	-3.5038	-0.4916	-0.1319	-0.0658	0.2733	4.3649
Total	-3.5038	-0.4988	-0.1351	-0.0744	0.2664	4.3649

toward a more strongly skewed underlying distribution would imply a concentration effect “around the flag” conditional on the first two moments remaining the same. Furthermore, it implies that bookmakers, to some extent, copy each other because of uncertainty about how to react and not because of new information. Otherwise, all four moments would remain the same, or at least one of the first two would change.

Figure 9 displays the results for the effect of newly announced infection on betting odds. One can easily see that there is no effect on the respective opponent but a significantly negative effect on γ_m for the second half of the sample. In Figure 10, the dynamic event study for the effect of a new (late) infection on betting odds' skewness. One can see a large effect for the first matchday after an infection, but only significant at the 10% level.

Effect on the kurtosis of the odds: The last moment we consider is the kurtosis. It looks at the “peakedness” of a distribution and, likewise, at the extreme tails. By that, it corresponds to the variance to a certain extent but adds further information. Table 2 presents summary statistics of this moment. Early infections have a kurtosis highly similar to the normal distribution ($\kappa_{\text{norm}} = 3$). Odds distributions for late infections as well as without infections are leptokurtic. They have more extreme outliers in odds and a higher frequency.

Looking at the response of the kurtosis to a newly announced infection, we do not observe a measurable reaction. As one can draw from Figure 11, the coefficients for baseline, late, and early infections are insignificant for the affected and the opponent team. As before, we look at the dynamic event study for late infections. We observe a flat pre-trend and a significant negative effect for the second match after a new infection. Put differently, the tails get thinner for the second match, and the peak around the mean gets larger.

Variation in the density functions: The previous findings are all captured in the differences of the probability density functions (PDF) as displayed in Figure 12. The blue line plots the density function for the matches prior to a

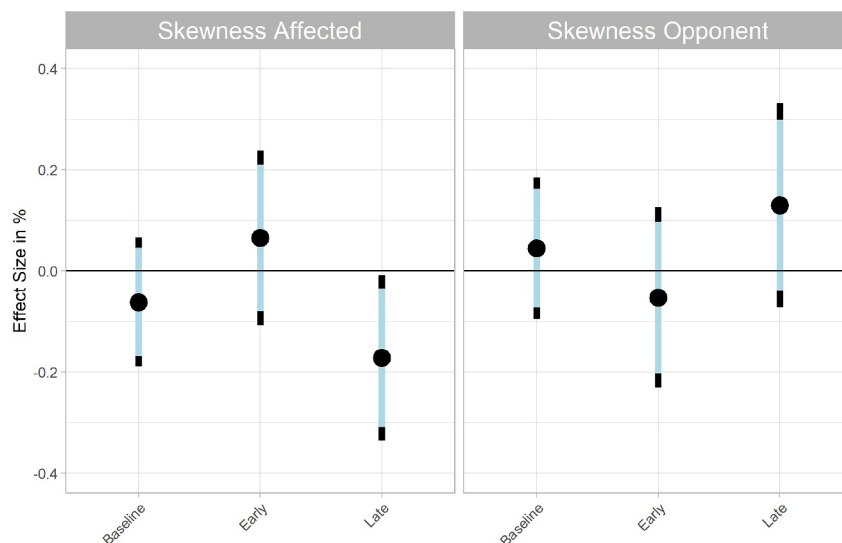


FIGURE 9 Dynamic Effect of a new absence on the skewness of betting odds. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the skewness (in levels) of the matchday odds for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).




FIGURE 10 Dynamic Effect on the skewness of the odds for a win for late infections. This figure plots the OLS coefficients β_T of the DiD regression following Equation (2) for the skewness (in levels) of the matchday odds for teams with new infections (LHS) within the subsample of late infections. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

new absence announcement. The red and the green lines show the first match after a new announcement. For both “post announcement” functions, one can see a higher level of dispersion and a strong shift in the mean for the PDF of late infections. Furthermore, a shift toward a more negatively skewed odds distribution is revealed, especially for the late infection announcements (green line). The visual changes in the actual distributions of odds back the regression analyses and highlight the differences between the market reaction to early and late absence announcements.

TABLE 2 Summary statistics on the kurtosis of betting odds.

Sample	Minimum	1 th quartile	Median	Mean	3 th quartile	Maximum
Early inf.	1.515	2.393	2.925	3.086	3.418	6.096
Late inf.	1.636	2.444	3.060	3.537	3.655	21.529
No inf.	1.433	2.391	2.927	3.414	3.766	22.128
Total	1.433	2.392	2.930	3.409	3.749	22.128

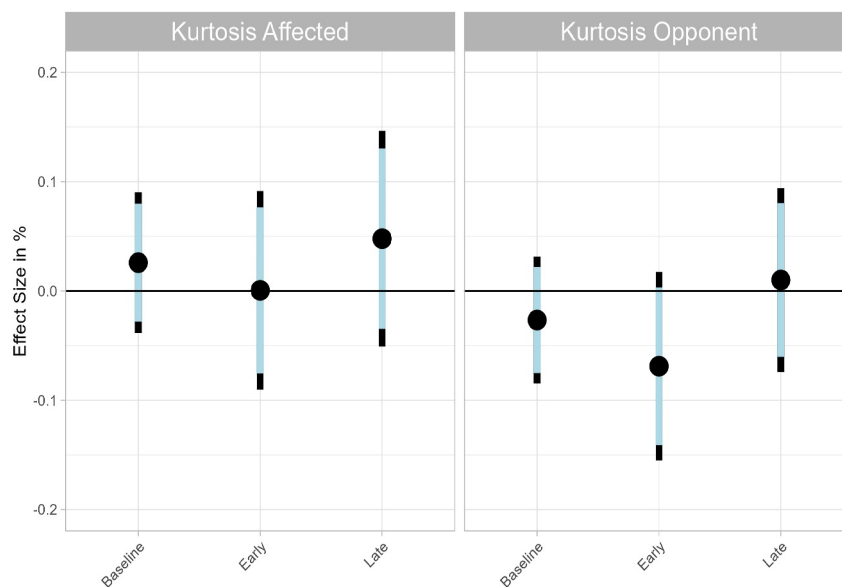


FIGURE 11 Effect of a new absence on the kurtosis of betting odds. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the kurtosis (in levels) of the matchday odds for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

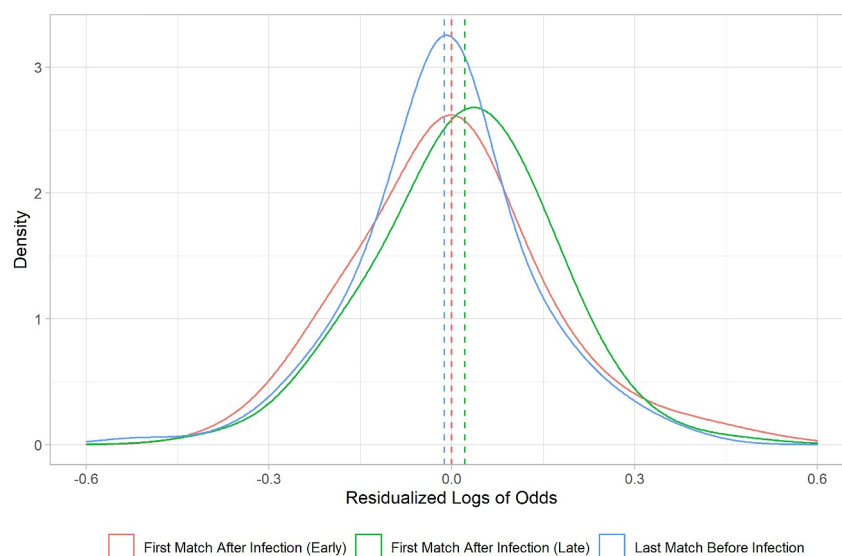


FIGURE 12 Probability density functions (PDF) of the odds distributions. This figure plots the probability density function of the residualized logs of odds. We residualize odds by regressing logged odds on team-season, opponent-season, and matchday-season fixed effects. The blue line plots the distribution of the last match before the announcement of a new infection. The red and green lines show the distribution for the first match after an absence announcement for early (red) and late infections (green), separated at the median infection date.

How can the movement of these statistical moments be interpreted in economic terms? Without a new absence announcement, we observe an approximately normal distribution of the residualized logs of the betting odds. In particular, we identify a shift in the first three moments for late infections (green line in Figure 12). The increase in the mean odds hints at the aggregated perception that the probability of winning a match following a new absence is lower than before. However, the changes in standard deviation and skewness emphasize that there is not just a shift to the right of the distribution.

We see bigger swings in the guesses of the bookmakers captured by the dispersion. Furthermore, the increase in the skewness (in absolute terms) is evidence that some bookmakers shift toward higher odds, which are gathered a bit closer around the mode. Still, fewer guesses are at the odds distribution's upper end. The long tail on the left side of the distribution also hints at bookmakers sticking to their odds.¹⁷

In addition, we run quantile regressions to obtain more information on which area of the odds distribution the treatment effect exerts its most substantial influence. As Table A4 in the Appendix illustrates, the effect on the first two moments is only significant for the first five and highest for the two lowest deciles. Thus, the change happens in particular for low odds, which, in turn, implies that bookmakers raise the odds of the favorite teams of a match, plausibly considering their chances to win to be lower after the announcement of a newly missing player. Therefore, the market's mispricing mostly appears for favorites, even though these are the clubs for which arguably the most information is available.

5.2 | Effect heterogeneity

Effect of Announcement Timing: When considering price adjustments, the temporal distance to the event is crucial. An announcement of a new absence far ahead of the next match not only allows market participants to spend a sufficient amount of time reevaluating their predictions. It also allows them to wait for and observe the reactions of other betting companies. Hence, the effect detected in the previous regressions will most likely be the largest (if existent at all) for announcements close to matchdays as imperfect information might be “traded away” over time. This is the case, as Figure 13 shows. For new notices within a range of 1-3 and 4-6 days ahead of a match, we can identify a significant change on the 5% level in both the mean and the standard deviation of the odds. For announcements seven up to 15 days ahead of a match, we do not find any impact of a new absence on these two moments.¹⁸ Overall, only announcements that would require an urgent response (if at all) trigger a substantial market reaction.

Player and Infection Event Heterogeneity: The effect of a missing player on the strength of a team crucially depends on the player's role within their squad. Imagine as two extreme examples the impact of either a replacement goalkeeper, which is likely to have no effect. In contrast, absent stars may have a much stronger impact on team performance. We use the individual share of matches played and relative market value to approximate a player's impact and relative role in their team (using *Transfermarkt* data).¹⁹ The upper two panels of Figure 14 show these




FIGURE 13 Effect decomposition by time between announcement and matchday. This figure plots the OLS coefficients for the first two moments decomposed by the timing of the announcement relative to the distance to the following matchday. Left panel: Mean; right panel: Standard deviation. Standard errors: Heteroskedasticity-robust and clustered at the match level. 90% (blue lines) and 95% confidence bands (black extensions) are applied.

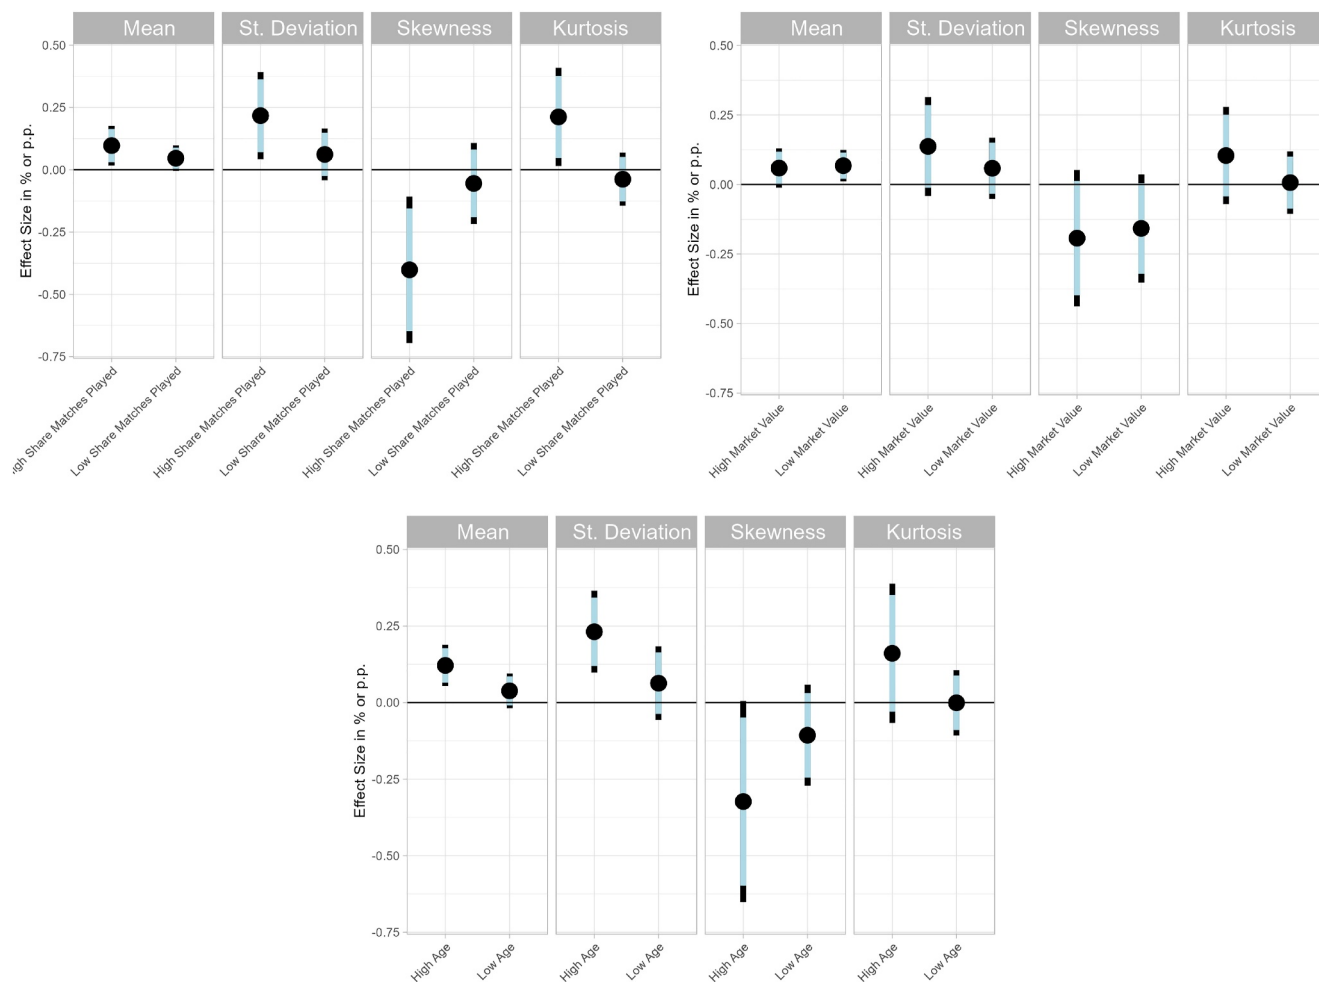


FIGURE 14 Effect decomposition by player relevance (top panels) and age (bottom panel). These figures plot the OLS coefficients for late infections of the DiD regression following Equation (1) for the first four moments of the matchday odds for teams with new infections. The panels differentiate between the number of matches played, market values, and player age. We compare the top tercile (high share matches played or age) with the bottom and middle tercile (low share matches played or age) of all infected players for both named dimensions of heterogeneity. For market values, “high market value” implies a player has a higher market value than the average market value of his team in the period before his infection. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

heterogeneities. A newly announced absence of high-impact players (i.e., a high share of matches played or high market value) significantly affects betting odds for late infections, while we do not detect any robust reaction of the odds’ moments to the absences of low-impact players.

Further, the lower panel of Figure 14 presents heterogeneity by the players’ age. Fischer et al. (2022) show that absent players’ performance after a COVID-19 infection is affected the most for older players. As for some infections, the first match post-infection is already after the mandatory quarantine period, so the health effects on players are important to account for. Indeed, betting odds are affected the most for older infected players.

Lastly, we study whether the effect of infection events varies with the number of infections. Figure 15 shows the effect of the treatment on the moments of the odds distribution. No difference between single- and multiple-infection events is evident.

5.3 | Hypothesis evaluation

In summary, a diffuse picture emerges from the changes in the statistical moments. Table 3 shows that none of the four moments of the betting odds distributions change among early infections. Quite the opposite holds for late infections.

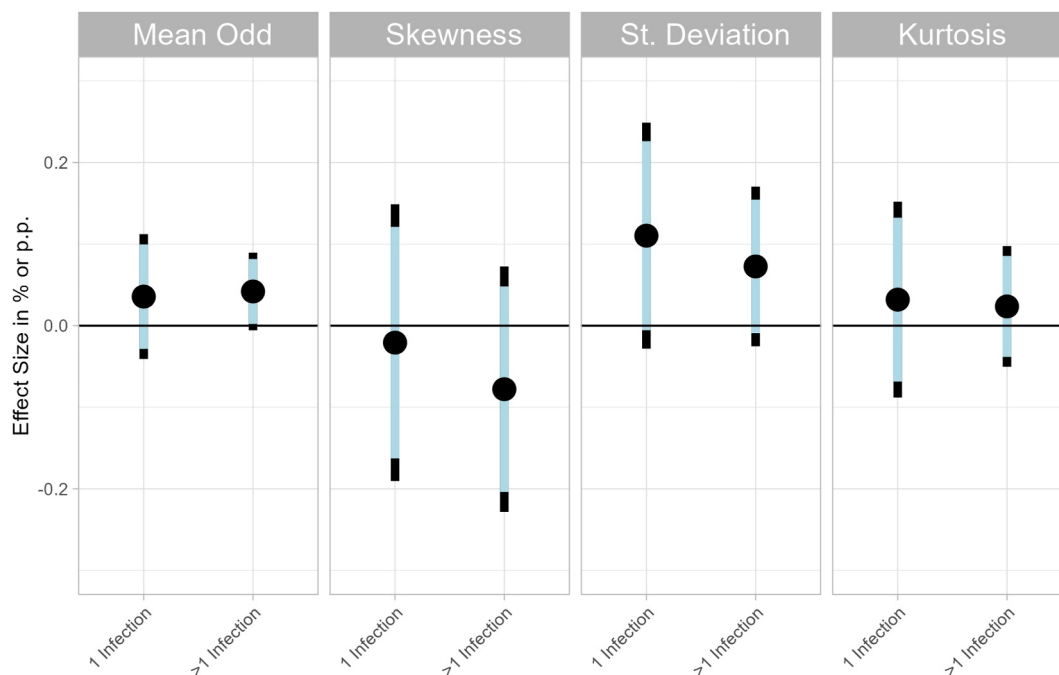


FIGURE 15 Effect decomposition by the number of new infections. This figure plots the OLS coefficient of the DiD regression for the first four moments of the matchday odds for teams with new infections. We estimate the interaction of the infection treatment with a dummy for single-infection events and a dummy for multiple-infection events. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

TABLE 3 Change in statistical moments of betting odds distributions.

Moment	Early infections	Late infections
Change in mean	–	✓
Change in standard deviation	–	✓
Change in skewness	–	✓
Change in kurtosis	–	–

Here, we observe a change in the first three moments and no variation only for the kurtosis. For the latter, there exists, at least, a slight movement toward a more platykurtic distribution for the second post-infection match as shown in the event study in Figure 16.

Regarding our proposed behavioral hypotheses, Table 4 shows that the only explanatory approach we can confirm is Hyp. 1 for the first half of the infections. The second half provides a more puzzling setting. While the change in the mean odds argues in favor of an educated adjustment of odds (i.e., Hyp. 2), the change in the standard deviation supports Hyp. 3, hence the continued existence or even increase of uncertainty. The change in skewness hints at some concentration, but it is challenged by the changes in the first two moments. It is supported by a slight change in the kurtosis for later infections, which hints, together with the skewness effect, toward “rally around the flag” behavior. Overall, the advent of new information at first leads to inertia, which could be considered *prima facie* as a learning process. After this initial shock, betting odds, as it were, frolic around, captured by $\Delta E_m[\xi_i] \neq 0$, $\Delta V_m[\xi_i] \neq 0$, and $\Delta |\gamma_m[\xi_i]| < 0$ at the same time. This stands in contrast with more common types of uncertainty in betting markets.

Deutscher et al. (2018) find inefficiencies in sports betting during the first matchdays of a season for recently promoted teams—a case with less information on team and player strengths than usual. It rapidly diminishes—other than the market distortions detected in our setting of new absence announcements. As previously stated, missing players did not affect the core outcomes of professional soccer. After some inevitable adjustment and belief updating, betting markets should have returned to no significant change in all four moments.

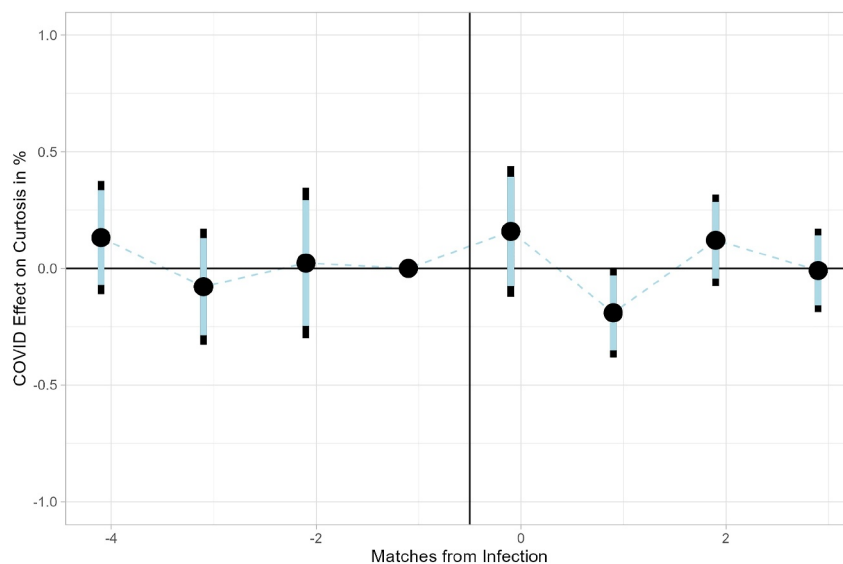


FIGURE 16 Dynamic Effect on the kurtosis of the odds for a win for late infections. This figure plots the OLS coefficients β_T of the DiD regression following Equation (2) for the kurtosis (in levels) of the matchday odds for teams with new infections (LHS) within the subsample of late infections. Standard errors: Heteroskedasticity-robust and clustered at the match level. 90% (blue lines) and 95% confidence bands (black extensions) are shown.

TABLE 4 Validation of stated behavioral hypotheses.

Hypothesis	$\Delta E_m[\xi_i]$	$\Delta V_m[\xi_i]$	$\Delta \gamma_m[\xi_i]$	$\Delta \kappa_m[\xi_i]$
Early infections				
1: Null effect	✓	✓	✓	✓
2: Educated adjustment	—	✓	✓	✓
3: Uneducated uncertainty	—	—	—	—
4: Rally round the flag	—	✓	—	—
Late infections				
1: Null effect	—	—	—	✓
2: Educated adjustment	✓	✓	—	(—)
3: Uneducated uncertainty	—	✓	—	—
4: Rally round the flag	—	—	✓	(✓)

At first glance, the observed overreaction might be explainable with the well-known representative heuristic (Tversky & Kahneman, 1974), which states that people overweight recent and salient information while neglecting the base rate probabilities. It occurs in financial markets (Daniel et al., 1998; De Bondt & Thaler, 1985) and empirical investigations hint especially at verdant agents being directed by that bias (Grether, 1980). Notwithstanding, bookmakers are highly experienced firms that possess a lot of past data and knowledge. Not only being heavily influenced by those well-established events of absence announcements but also being affected by late announcements appears preposterous. Nevertheless, we observe these movements in the data. It could hint that experienced agents also suffer from distortions in decision-making that overrule potentially sophisticated forecasting techniques. Likewise, it could also be caused by distorted beliefs of the bettors to which the bookmakers adjust their odds.

Considering our findings as they are, what can we draw from them? Even rational market participants need time to understand new information's value and implications. Looking at our dynamic results, we see evidence precisely for that. We identify significant changes in the first three moments, at least on the 10% level, for the first matchday after an announcement that disappears right after. While the reaction at this first post-infection match is somewhat erratic,

reverting to a null effect is in line with the benchmark that cannot find an effect of an additional missing player on match outcomes. However, this only holds for the second half of the sample. For the first half, there is neither a dynamic effect nor an effect on the static semi-elasticity.

Stepping back, we can arguably confirm the Hayek hypothesis *over time* for each individual infection (only shortlived reaction that returns to the benchmark suggestion), but not *across* infections. Why betting markets appear incapable of extrapolating from one player to another remains puzzling. In particular, it is surprising that all three moments change for late infections. Given that players missing for some reason is not a new phenomenon, we cannot rationalize why odds “frolic around.” In the same sense, it is puzzling that there is no reaction for early infections since quality deterioration should also be an issue among these infections, and uncertainty has risen. Thus, our case of betting markets raises doubts as to the extent to which the Hayek hypothesis holds for adequate incorporation of new information.

6 | DISCUSSION

Using the global COVID-19 pandemic as a catalyst for absenteeism on the pitch, we can identify substantial inertia and behavioral patterns that hardly confirm the strength of markets in fast information processing for the high-stakes environment of betting markets. Both bettors and bookmakers are interested in the fast processing of new information and evaluating whether it is relevant for match outcomes on which one can bet. Confronted with infections of elite athletes in Europe, the odds of bookmakers tend not to react at all to new infections in the beginning. Only after a long period of inertia do we identify a puzzling mixture of a higher mean, higher dispersion, and a more left-skewed distribution of odds (and implied probabilities of a win) that do not follow a clear rationale. It violates the hypothesis of semi-strong market inefficiency.

While we cannot identify a reasonable behavioral pattern across new announcements, we might see some adjustment or learning for the time after a specific announcement. Among those matches in which betting odds changed in the first place, they seem to revert to status *ex-ante* after one match of “confusion.” However, this could be driven by the return to the pitch of many absent players after one missed match. While learning would back the Hayek hypothesis of quick information transmission within markets and rational learning, it remains an open question why we find an effect only for a subsample later in time—and, of course, why odds react in such a way in the first place.

Exploiting COVID-19 as a high-frequent information shock for betting markets has demonstrated that price adjustments do not necessarily reflect the *ex-post* benchmark. The actual response to the new (and apparently meaningless) information seems to consist of inertia combined with lagged blind action: Even though we cannot rule out alternative explanations for the null effect among early infections, an explanation that argues in favor of rational behavior in response to this first half of the sample lacks a sufficient explanation, why market participants alter their behavior in the second half of the sample. One has to remember that changes in betting odds can be either caused by the bookmakers mispricing new information or the bettors falsely updating their individual expectations on match outcomes.

Our analysis raises concerns to which extent the strength of market environments—as captured by the Hayek hypothesis (Smith, 1982)—still holds for the advent of new information. The data reflect the unique example of a shock with hitherto unknown consequences and the respective external evaluation. Furthermore, the analysis provides much more substantial evidence as we use agents with a high commitment to correct outcomes and a large variety of bookmakers, so we are confident we can rule out random effects caused by a selective sample of data. Eventually, many global and industry-specific challenges require belief updating sufficiently fast. Overly long reaction times are a significant obstacle to this. As cited earlier, Buchanan (1964, p. 218) states that “a market is not competitive by assumption or by construction. A market becomes competitive.” At least for our European sports betting market sample, we find the opposite. It is an important finding and a puzzle at the same time. Future research would need to disentangle the interdependencies of bettor and bookmaker behavior in more detail and how they interact.

ACKNOWLEDGMENTS

We thank Christian Deutscher, Justus Haucap, Simon Martin, and J. James Reade as well as participants of the 11th Football & Finance Conference at Mercator School of Management 2024, the Leipzig Colloquium on the Foundations of the Market Order 2023, the ESEA Annual Conference 2022, the IX. MathSports Conference 2022, and the ROSES research seminar for helpful feedback. We thank *transfermarkt.de* for providing us with the data on players' market

values. We gratefully acknowledge the hospitality of the University of Reading's Department of Economics during our research visit and funding from the German Research Foundation (DFG) — 235577387/GRK1974.

Open Access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in openIPSCR at <https://doi.org/10.3886/E207770V5>, reference number 207770 (Fischer & Schmal, 2024).

ORCID

W. Benedikt Schmal  <https://orcid.org/0000-0003-2400-2468>

ENDNOTES

- ¹ See, for example, Arrow et al. (2008), Hanson et al. (2006), Kremer et al. (2014), Snowberg et al. (2013), and Wolfers and Zitzewitz (2004).
- ² European Online Gambling Key Figures 2021 Edition, <https://www.egba.eu/uploads/2021/12/European-Online-Gambling-Key-Figures-2021-Edition.pdf>, including EU-27 countries + UK.
- ³ Even though betting odds are correlated across bookmakers, this is not an obstacle to our analysis as we specifically study the reaction to the new information that may be more heterogeneous.
- ⁴ The hypothesis has been proven to hold in various experimental settings (Bulte et al., 2017; Davis & Williams, 1991; Shachat & Zhang, 2017).
- ⁵ The website <https://www.oddsportal.com/> is a major archive of odds in various sports and has already provided data for research (Gomez-Gonzalez et al., 2019; Qasim et al., 2020; Reade et al., 2021). Table A1 in the Appendix lists all bookmakers included.
- ⁶ In theory, we could expand our analysis to further leagues. However, the causality of our econometric design relies crucially on a clean identification of infected and non-infected players. Fischer et al. (2022) provide that only for Italy and Germany, and no other studies have gathered infections in such great detail.
- ⁷ Consider, for example, teams *A* and *B* playing. The odds for a win of the home team (1) is for team *A* the odds for its win. In contrast, for *B*, it is the odds for a win of the opponent team (2). Of course, the whole set-up, to some extent, faces the “1 – *p*-problem” (Schmal, 2023) as a change in the likelihood of, for example, winning affects the other two odds for a match most likely in the opposite direction.
- ⁸ Figure A1 in the Appendix provides an overview of the size and frequency of these odds.
- ⁹ 91 matches out of the 131 matches (71%) in total are single infections. Hence, distinguishing between single and multiple infections would be interesting, but statistical power is lacking.
- ¹⁰ Our main regression results use the whole panel, that is, we have a varying number of bookmakers because, for some matchdays, odds from some bookmakers are missing. It is unclear whether this is caused by missing values on *oddsportal* or by some bookmakers not offering odds for all matches. We conduct an additional analysis using a balanced panel that only includes bookmakers with a full set of odds. Figure A2 in the appendix shows that the results are qualitatively equivalent.
- ¹¹ The regression results for a split into quartiles and quintiles in Figure A9 in the Appendix generally confirm our findings.
- ¹² Additionally, we look at the coefficient of variation (CVR) because the size of the standard deviation depends on the expected value of the odds. In contrast, the coefficient of variation is stationary for log-normal distributed values. Given the distribution of odds, it is reasonable to assume log-normality. We construct the raw coefficient of variation as $\ln[CVR]_{jm} = \sqrt{e^{\hat{\sigma}_{\ln|jm}^2} + 1}$, where $\hat{\sigma}_{\ln|jm}^n$ is the empirical standard deviation of $\ln(\xi_{jm}) \forall i$. Figure A6 in the Appendix presents the results.
- ¹³ This setup has been questioned in a recently developing literature, which identified problems with using event studies to estimate treatment effects from a staggered roll-out. In particular, comparing already-treated to not-yet-treated observations can distort estimated coefficients. Admittedly, several correction estimators are proposed in the literature (see, e.g., the survey paper by De Chaisemartin & d'Haultfoeuille, 2023). However, these models are designed for absorbing treatment, that is, each individual is treated only once. In our setting, clubs face new infections multiple times. Applying the correction only to clubs with at most one infection event or to clubs' first infection events is inappropriate in our setting since we are particularly interested in the effect heterogeneity over time.
- ¹⁴ The prediction error is the difference between a team's outcome and the implied probability of this outcome elicited from the betting odds.
- ¹⁵ As a robustness check, we have also computed the raw coefficient of variation. The results align with those for the standard deviation and can be found in Figure F6 in the Appendix.
- ¹⁶ Comparing the p90-p10 quantile distance, we also find an increase that hints at growth in dispersion (see Figure A7 in the Appendix). This measure is well-established in work on labor markets and inequality (Gottschalk & Smeeding, 1997).
- ¹⁷ Alas, we cannot disentangle this stickiness of those bookmakers who remain inert and those who do not adjust their odds because they correctly expect no changes in the likelihood of winning.
- ¹⁸ Figure A8 in the Appendix plots the decomposition by announcement timing for skewness. We do not find significant effects for the decomposition by temporal distance to the next matchday.

¹⁹ We compare a player's market value with the mean market value of his team before his infection.

REFERENCES

- Al-Ubaydli, O., Boettke, P. & Albrecht, B.C. (2022) Testing the Hayek hypothesis: recent theoretical and experimental evidence. *PLoS One*, 17(7), 1–25. Available from: <https://doi.org/10.1371/journal.pone.0270489>
- Alderfer, C.P. & Bierman, H. (1970) Choices with risk: beyond the mean and variance. *Journal of Business*, 43(3), 341–353. Available from: <https://doi.org/10.1086/295286>
- Angelini, G. & De Angelis, L. (2019) Efficiency of online football betting markets. *International Journal of Forecasting*, 35(2), 712–721. Available from: <https://doi.org/10.1016/j.ijforecast.2018.07.008>
- Angelini, G., De Angelis, L. & Singleton, C. (2022) Informational efficiency and behaviour within in-play prediction markets. *International Journal of Forecasting*, 38(1), 282–299. Available from: <https://doi.org/10.1016/j.ijforecast.2021.05.012>
- Arrow, K.J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J.O., et al. (2008) The promise of prediction markets. *Science*, 320(5878), 877–878. Available from: <https://doi.org/10.1126/science.1157679>
- Avery, C. & Chevalier, J. (1999) Identifying investor sentiment from price paths: the case of football betting. *Journal of Business*, 72(4), 493–521. Available from: <https://doi.org/10.1086/209625>
- Balanda, K.P. & MacGillivray, H.L. (1988) Kurtosis: a critical review. *The American Statistician*, 42(2), 111–119. Available from: <https://doi.org/10.1080/00031305.1988.10475539>
- Berg, J.E., Nelson, F.D. & Rietz, T.A. (2008) Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2), 285–300. Available from: <https://doi.org/10.1016/j.ijforecast.2008.03.007>
- Brown, A. (2014) Information processing constraints and asset mispricing. *The Economic Journal*, 124(575), 245–268. Available from: <https://doi.org/10.1111/eoj.12057>
- Brown, A. & Yang, F. (2016) Limited cognition and clustered asset prices: evidence from betting markets. *Journal of Financial Markets*, 29, 27–46. Available from: <https://doi.org/10.1016/j.finmar.2015.10.003>
- Brown, A., Reade, J.J. & Williams, L.V. (2019) When are prediction market prices most informative? *International Journal of Forecasting*, 35(1), 420–428. Special Section: Supply Chain Forecasting. Available from: <https://doi.org/10.1016/j.ijforecast.2018.05.005>
- Buchanan, J.M. (1964) What should economists do? *Southern Economic Journal*, 30(3), 213–222. Available from: <https://doi.org/10.2307/1055931>
- Bulte, E., Kontoleon, A., List, J., Turley, T. & Voors, M. (2017) From personalized exchange towards anonymous trade: a field experiment on the workings of the invisible hand. *Journal of Economic Behavior & Organization*, 133, 313–330. Available from: <https://doi.org/10.1016/j.jebo.2016.10.014>
- Choi, D. & Hui, S.K. (2014) The role of surprise: understanding overreaction and underreaction to unanticipated events using in-play soccer betting market. *Journal of Economic Behavior & Organization*, 107, 614–629. Available from: <https://doi.org/10.1016/j.jebo.2014.02.009>
- Croxson, K. & Reade, J.J. (2014) Information and efficiency: goal arrival in soccer betting. *The Economic Journal*, 124(575), 62–91. Available from: <https://doi.org/10.1111/eoj.12033>
- Daniel, K., Hirshleifer, D. & Subrahmanyam, A. (1998) Investor psychology and security market under- and overreactions. *The Journal of Finance*, 53(6), 1839–1885. Available from: <https://doi.org/10.1111/0022-1082.00077>
- Davis, D. & Williams, A. (1991) The Hayek hypothesis in experimental auctions: institutional effects and market power. *Economic Inquiry*, 29(2), 261–274. Available from: <https://doi.org/10.1111/j.1465-7295.1991.tb01270.x>
- De Bondt, W.F.M. & Thaler, R. (1985) Does the stock market overreact? *The Journal of Finance*, 40(3), 793–805. Available from: <https://doi.org/10.2307/2327804>
- De Chaisemartin, C. & d'Haultfoeulle, X. (2023) Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. *The Econometrics Journal*, 26(3), C1–C30. Available from: <https://doi.org/10.1093/ectj/utac017>
- Deck, C. & Porter, D. (2013) Prediction markets in the laboratory. *Journal of Economic Surveys*, 27(3), 589–603. Available from: <https://doi.org/10.1111/joes.12015>
- Dertwinkel-Kalt, M. & Köster, M. (2019) Salience and skewness preferences. *Journal of the European Economic Association*, 18(5), 2057–2107. Available from: <https://doi.org/10.1093/jeea/jvz035>
- Deutscher, C., Frick, B. & Ötting, M. (2018) Betting market inefficiencies are short-lived in German professional football. *Applied Economics*, 50(30), 3240–3246. Available from: <https://doi.org/10.1080/00036846.2017.1418082>
- Doane, D.P. & Seward, L.E. (2011) Measuring skewness: a forgotten statistic? *Journal of Statistics Education*, 19(2). Available from: <https://doi.org/10.1080/10691898.2011.11889611>
- Ebert, S. (2015) On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, 112, 85–97. Available from: <https://doi.org/10.1016/j.jebo.2015.01.003>
- Fama, E.F. (1970) Efficient capital markets: a review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417. Available from: <https://doi.org/10.1111/j.1540-6261.1970.tb00518.x>
- Feddersen, A., Humphreys, B.R. & Soebbing, B.P. (2017) Sentiment bias and asset prices: evidence from sports betting markets and social media. *Economic Inquiry*, 55(2), 1119–1129. Available from: <https://doi.org/10.1111/ecin.12404>
- Fischer, K. & Haucap, J. (2022) *Home advantage in professional soccer and betting market efficiency: the role of spectator crowds*. Kyklos. forthcoming.
- Fischer, K. & Schmal, W.B. (2024) ECIN replication package for “Pricing in response to new information: the case of betting markets.” *openICPSR*. Available from: <https://doi.org/10.3886/E207770V5>

- Fischer, K., Reade, J.J. & Schmal, W.B. (2022) What cannot be cured must be endured: the long-lasting effect of a COVID-19 infection on workplace productivity. *Labour Economics*, 79, 102281. Available from: <https://doi.org/10.1016/j.labeco.2022.102281>
- Golec, J. & Tamarkin, M. (1998) Bettors love skewness, not risk, at the horse track. *Journal of Political Economy*, 106(1), 205–225. Available from: <https://doi.org/10.1086/250007>
- Gomez-Gonzalez, C., del Corral, J., Jewell, R.T., Garcia-Unanue, J. & Nesseler, C. (2019) A prospective analysis of competitive balance levels in major league soccer. *Review of Industrial Organization*, 54(1), 175–190. Available from: <https://doi.org/10.1007/s11151-018-9667-3>
- Gottschalk, P. & Smeeding, T.M. (1997) Cross-national comparisons of earnings and income inequality. *Journal of Economic Literature*, 35(2), 633–687.
- Grether, D.M. (1980) Bayes rule as a descriptive model: the representativeness heuristic. *Quarterly Journal of Economics*, 95(3), 537–557. Available from: <https://doi.org/10.2307/1885092>
- Grossman, S. (1976) On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2), 573–585. Available from: <https://doi.org/10.2307/2326627>
- Grossman, S. (1978) Further results on the informational efficiency of competitive stock markets. *Journal of Economic Theory*, 18(1), 81–101. Available from: [https://doi.org/10.1016/0022-0531\(78\)90043-1](https://doi.org/10.1016/0022-0531(78)90043-1)
- Hanson, R., Oprea, R. & Porter, D. (2006) Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4), 449–459. Available from: <https://doi.org/10.1016/j.jebo.2004.09.011>
- Hayek, F. (1945) The use of knowledge in society. *The American Economic Review*, 35(4), 519–530.
- Hellwig, M.F. (1980) On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3), 477–498. Available from: [https://doi.org/10.1016/0022-0531\(80\)90056-3](https://doi.org/10.1016/0022-0531(80)90056-3)
- Hirshleifer, J. (1965) Investment decision under uncertainty: choice—theoretic approaches. *Quarterly Journal of Economics*, 79(4), 509–536. Available from: <https://doi.org/10.2307/1880650>
- Kremer, I., Mansour, Y. & Perry, M. (2014) Implementing the “wisdom of the crowd”. *Journal of Political Economy*, 122(5), 988–1012. Available from: <https://doi.org/10.1086/676597>
- Levitt, S.D. (2004) Why are gambling markets organised so differently from financial markets? *The Economic Journal*, 114(495), 223–246. Available from: <https://doi.org/10.1111/j.1468-0297.2004.00207.x>
- Mao, J.C. (1970) Survey of capital budgeting: theory and practice. *The Journal of Finance*, 25(2), 349–360. Available from: <https://doi.org/10.2307/2325481>
- Meyer, A. & Hundtofte, S. (2023) The longshot bias is a context effect. *Management Science*, 69(11), 6954–6968. Available from: <https://doi.org/10.1287/mnsc.2023.4684>
- Moskowitz, T.J. (2021) Asset pricing and sports betting. *The Journal of Finance*, 76(6), 3153–3209. Available from: <https://doi.org/10.1111/jofi.13082>
- Muth, J.F. (1961) Rational expectations and the theory of price movements. *Econometrica*, 29(3), 315–335. Available from: <https://doi.org/10.2307/1909635>
- Ötting, M., Deutscher, C., Singleton, C. & De Angelis, L. (2023) *Gambling on momentum in contests*. Working Paper.
- Qasim, M., Kibria, B.M.G., Mansson, K. & Sjölander, P. (2020) A new Poisson Liu regression estimator: method and application. *Journal of Applied Statistics*, 47(12), 2258–2271. Available from: <https://doi.org/10.1080/02664763.2019.1707485>
- Quandt, R.E. (1986) Betting and equilibrium. *Quarterly Journal of Economics*, 101(1), 201–207. Available from: <https://doi.org/10.2307/1884650>
- Ramirez, P., Reade, J.J. & Singleton, C. (2023) Betting on a buzz: mispricing and inefficiency in online sportsbooks. *International Journal of Forecasting*, 39(3), 1413–1423. Available from: <https://doi.org/10.1016/j.ijforecast.2022.07.011>
- Reade, J.J., Singleton, C. & Brown, A. (2021) Evaluating strange forecasts: the curious case of football match scorelines. *Scottish Journal of Political Economy*, 68(2), 261–285. Available from: <https://doi.org/10.1111/sjpe.12264>
- Schmal, W.B. (2023) Vice versa: the decoupling of content and topic heterogeneity in collusion research. *Journal of Economic Surveys*. forthcoming. Available from: <https://doi.org/10.1111/joes.12600>
- Shachat, J. & Zhang, Z. (2017) The Hayek hypothesis and long-run competitive equilibrium: an experimental investigation. *The Economic Journal*, 127(599), 199–228. Available from: <https://doi.org/10.1111/ecco.12249>
- Shin, H.S. (1993) Measuring the incidence of insider trading in a market for state-contingent claims. *The Economic Journal*, 103(420), 1141–1153. Available from: <https://doi.org/10.2307/2234240>
- Singh, A.K., Gewali, L.P. & Khatiwada, J. (2019) New measures of skewness of a probability distribution. *Open Journal of Statistics*, 9(5), 601–621. Available from: <https://doi.org/10.4236/ojs.2019.95039>
- Smith, V.L. (1982) Markets as economizers of information: experimental examination of the ‘Hayek hypothesis’. *Economic Inquiry*, 20(2), 165–179. Available from: <https://doi.org/10.1111/j.1465-7295.1982.tb01149.x>
- Snowberg, E., Wolfers, J. & Zitzewitz, E. (2013) Prediction markets for economic forecasting. In: *Handbook of Economic Forecasting*. Elsevier, 2, pp. 657–687. Available from: <https://doi.org/10.1016/b978-0-444-53683-9.00011-6>
- Thaler, R.H. & Ziemba, W.T. (1988) Anomalies: parimutuel betting markets: racetracks and lotteries. *The Journal of Economic Perspectives*, 2(2), 161–174. Available from: <https://doi.org/10.1257/jep.2.2.161>
- Tversky, A. & Kahneman, D. (1974) Judgment under uncertainty: heuristics and biases: biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, 185(4157), 1124–1131. Available from: <https://doi.org/10.1126/science.185.4157.1124>
- Waltz, K. (1967) *Electoral punishment and foreign policy crises*. Domestic Sources of Foreign Policy, pp. 263–293.

- Whelan, K. (2024) Risk aversion and favourite–longshot bias in a competitive fixed-odds betting market. *Economica*, 91(361), 188–209. Available from: <https://doi.org/10.1111/ecca.12500>
- Whelan, K. & Hegarty, T. (2023) Calculating the bookmaker’s margin why bets lose more on average than you are warned. Available at SSRN 4374924.
- Williams, L.V. & Reade, J.J. (2016) Prediction markets, social media and information efficiency. *Kyklos*, 69(3), 518–556. Available from: <https://doi.org/10.1111/kykl.12119>
- Winkelmann, D., Deutscher, C. & Ötting, M. (2021) Bookmakers’ mispricing of the disappeared home advantage in the German Bundesliga after the COVID-19 break. *Applied Economics*, 53(26), 3054–3064. Available from: <https://doi.org/10.1080/00036846.2021.1873234>
- Winkelmann, D., Ötting, M., Deutscher, C. & Makarewicz, T. (2024) Are betting markets inefficient? Evidence from simulations and real data. *Journal of Sports Economics*, 25(1), 54–97. Available from: <https://doi.org/10.117/15270025231204997>
- Wolfers, J. & Zitzewitz, E. (2004) Prediction markets. *The Journal of Economic Perspectives*, 18(2), 107–126. Available from: <https://doi.org/10.1257/0895330041371321>

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Fischer, K. & Schmal, W.B. (2025) Pricing in response to new information: the case of betting markets. *Economic Inquiry*, 63(1), 236–264. Available from: <https://doi.org/10.1111/ecin.13258>

APPENDIX A

TABLE A1 List of bookmakers.

Bookmaker	#Matches covered
10Bet	1372
10 × 10bet	214
188BET	1372
1 x Bet	1372
888sport	1372
Bet365	1372
Bet-at-home	1372
BetInAsia	461
BetVictor	1372
Betfair	1372
Betfair exchange	1246
Betfred	1372
Betsafe	1372
Betsson	1372
Betway	1372
BoyleSports	1372
Bwin	1372
ComeOn	1372
Coolbet	1372
Curebet	135
Dafabet	1372
GGBET	1221

TABLE A1 (Continued)

Bookmaker	#Matches covered
Interwetten	1363
Lasbet	135
Leonbets	1372
Marathonbet	1372
NordicBet	1372
Parimatch	1371
Pinnacle	1372
Unibet	1372
William hill	1372
Youwin	1356
Total	39,058

Note: Betfair exchange is an exception in that sense that here players bet against each other instead of betting against the bookmaker. “Betfair exchange” only provides the platform to offer and match bets. Hence, it is user generated betting. Total number of matches in the dataset: 1372.

TABLE A2 Random assignment of infection events across team-match observations.

Dependent variable	$T_{tm} = 1$
Team-match specific outcomes	
$\mathbb{1}[\text{Italian League}]_{tm}$	0.060 (0.047)
$\mathbb{1}[\text{Home}]_{tm}$	-0.051 (0.045)
$T_{t,m-1}$	0.062 (0.039)
$\text{Points}_{t,m-1}$	0.117 (0.112)
Lagged betting market outcomes	
Mean margin $_{t,m-1}$	0.001 (0.015)
Mean odd $_{t,m-1}$	-0.138 (0.278)
SD odd $_{t,m-1}$	0.018 (0.040)
Skew odd $_{t,m-1}$	0.028 (0.071)

Note: This table gives regression results of the treatment variable T_{tm} on a variety of match- or team-specific information. The regressions are conditional on matchday-season fixed effects to account for, for example, changes in the likelihood of infections over time.

TABLE A3 Regression results efficiency test as in Angelini and De Angelis (2019).

$\epsilon_{tm} = 1[\text{Win}_{tm}] - p_{tm}$		
With p_{tm} implied probability of outcome		
Constant	−0.001 (0.006)	−0.001 (0.006)
Treat	−0.009 (0.040)	
Treat × early		−0.009 (0.056)
Treat × late		−0.008 (0.054)
N	2744	2744

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

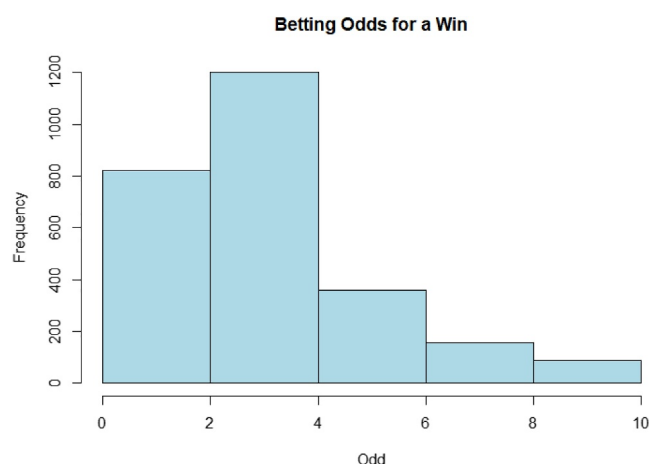


FIGURE A1 Histogram of mean betting odds for a win. Trimmed at mean odd of 10 (95th percentile).




FIGURE A2 Regression estimates first three moments—balanced panel. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the mean of the matchday odds for teams with new infections (LHS of each panel) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions). Dataset: *Balanced panel* only using bookmakers with odds for every match in the sample.




FIGURE A3 Histogram of matches with previous COVID-19 infections. Distribution of matches affected by new COVID-19 infections of players across time. Total number of matches with ≥ 1 players previously infected: 131.

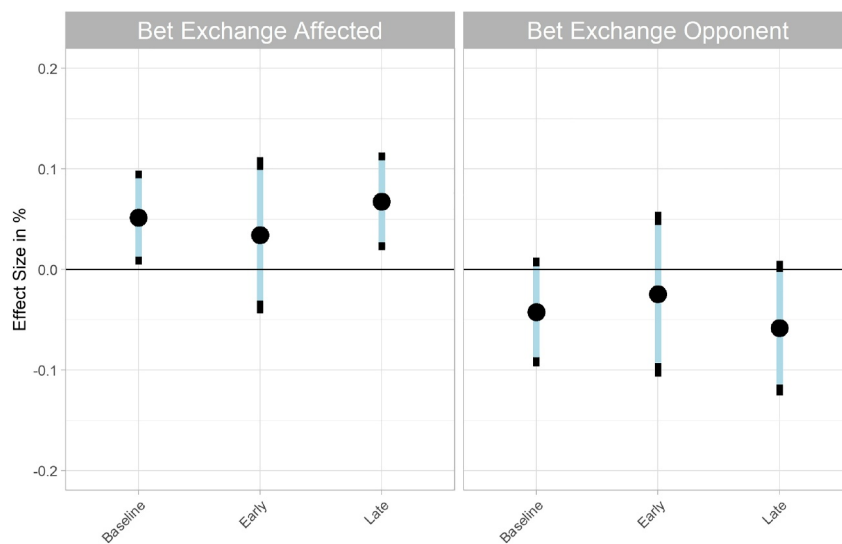


FIGURE A4 Odds comparison on “betfair exchange.” This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the mean of the matchday odds for teams with new infections (LHS) and their opponents (RHS) only looking at “betfair exchange” odds. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

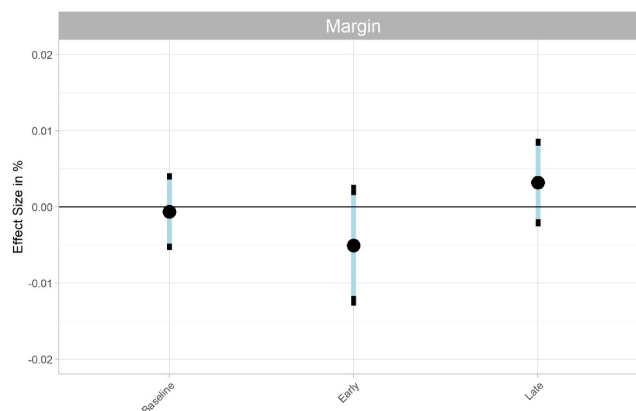


FIGURE A5 Effect of new absences on bookmaker margins. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for bookmaker margins (as defined by Shin, 1993) for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

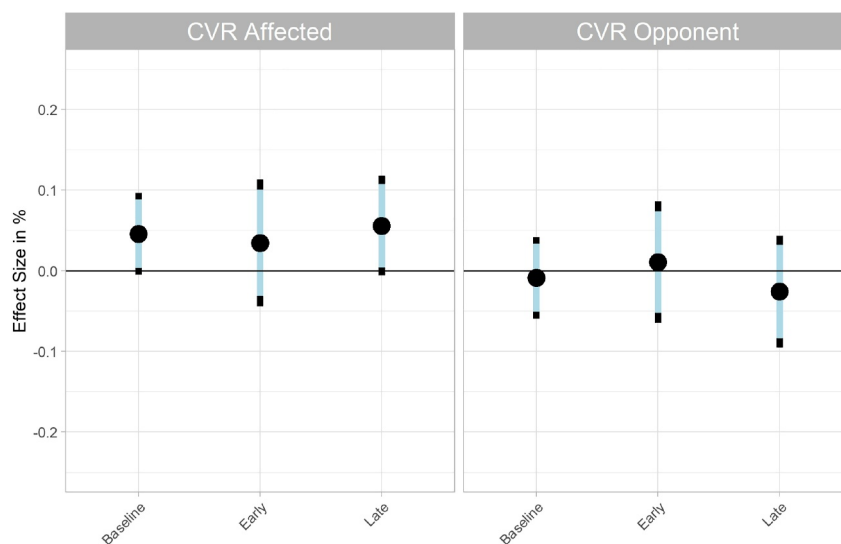


FIGURE A6 Effect on the coefficient of variation (CVR) of a win for a team with a new absence and its opponent. This figure plots the OLS coefficients β_T of the DiD regression following Equation (1) for the mean of the matchday odds for teams with new infections (LHS) and their opponents (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions). Dependent Variable: Raw CVR as defined in Section 3.

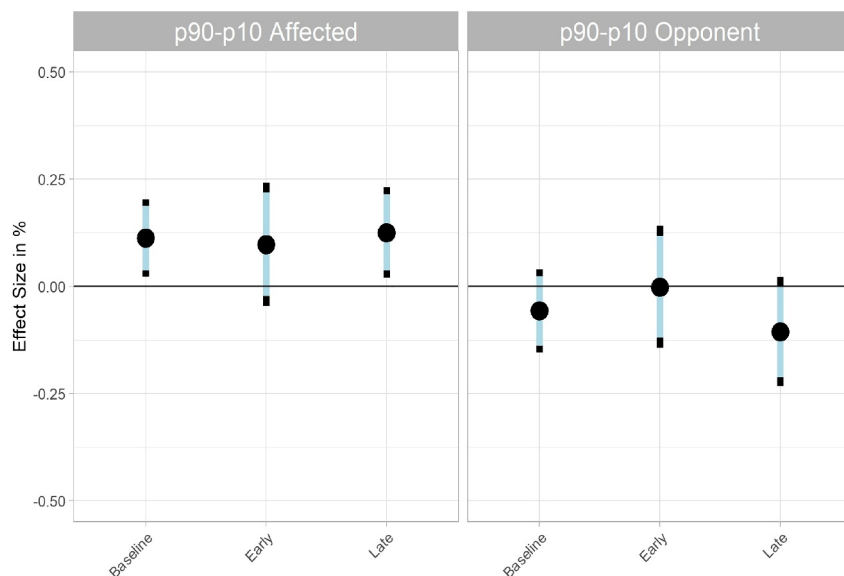


FIGURE A7 Effect on the dispersion of the odds: p90-p10 quantile distance. This figure plots the OLS coefficients of the DiD regression following Equation (1) for the p90-p10 quantile measure as alternative measure of dispersion. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).

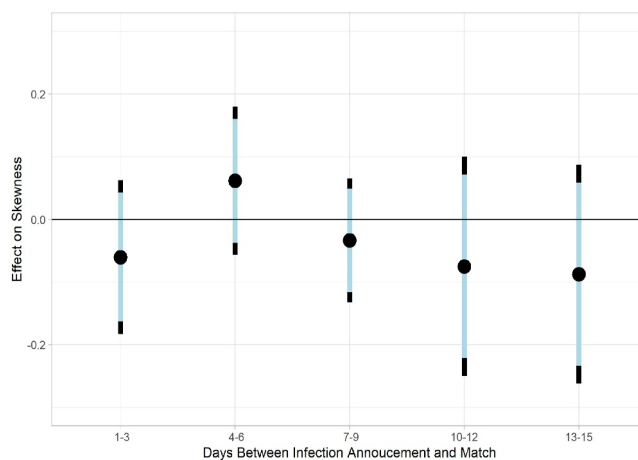


FIGURE A8 Skewness effect decomposition by time between announcement and matchday. This figure plots the OLS coefficients for the third moment decomposed by the timing of the announcement relative to the distance to the following matchday. Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (blue lines) and 95% (black extensions).




FIGURE A9 Regression estimates for separation of infections by quartiles and quintiles. This figure plots the OLS coefficients of the DiD regression following Equation (1) but not separated at the median infection but divided into quartiles (LHS) and quintiles (RHS). Standard errors: Heteroskedasticity-robust and clustered at the match level. Confidence bands: 90% (colored lines) and 95% (black extensions).

TABLE A4 Quantile regression on residualized first two moments.

	Residualized ln(Mean odd)		Residualized ln(St. Dev.)	
	$T_{tm} \times \mathbb{1}[\text{Early}]$	$T_{tm} \times \mathbb{1}[\text{Late}]$	$T_{tm} \times \mathbb{1}[\text{Early}]$	$T_{tm} \times \mathbb{1}[\text{Late}]$
$\tau = 0.1$	0.007 (0.050)	0.189*** (0.048)	-0.004 (0.128)	0.275*** (0.064)
$\tau = 0.2$	-0.003 (0.039)	0.115*** (0.020)	0.046 (0.077)	0.182*** (0.055)
$\tau = 0.3$	0.019 (0.055)	0.102*** (0.026)	0.091 (0.099)	0.112* (0.065)
$\tau = 0.4$	0.039 (0.034)	0.061** (0.029)	0.074 (0.087)	0.101 (0.068)
$\tau = 0.5$	0.011 (0.031)	0.061** (0.028)	0.134 (0.094)	0.119** (0.058)
$\tau = 0.6$	0.017 (0.034)	0.040 (0.023)	0.089 (0.055)	0.072 (0.052)
$\tau = 0.7$	0.010 (0.029)	0.027 (0.025)	0.043 (0.038)	0.055 (0.063)
$\tau = 0.8$	0.010 (0.049)	-0.023 (0.039)	-0.056 (0.054)	-0.006 (0.059)
$\tau = 0.9$	-0.021 (0.060)	-0.022 (0.032)	-0.094 (0.162)	-0.034 (0.114)

Note: We residualize the logged outcomes in a linear regression on team-season, opponent-season and matchday-season fixed effects. We then run a quantile regression on the two independent treatment variables for all deciles. Bootstrapped standard errors in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.