



**University of
Zurich^{UZH}**

Department of Business Administration

**UZH Business Working Paper Series
(ISSN 2296-0422)**

Working Paper No. 368

**Insider trading and price efficiency:
Evidence from a betting exchange**

Paolo Bizzozero, Raphael Flepp, and Egon Franck

September 2017

University of Zurich, Plattenstrasse 14, CH-8032 Zurich,
<http://www.business.uzh.ch/forschung/wps.html>

UZH Business Working Paper Series

Contact Details

Paolo Bizzozero

paolo.bizzozero@business.uzh.ch

Raphael Flepp

raphael.flepp@business.uzh.ch

Egon Franck

egon.franck@business.uzh.ch

University of Zurich

Department of Business Administration

Affolternstrasse 56, CH-8050 Zurich, Switzerland

Tel.: +41 44 634 29 13

Insider Trading and Price Efficiency: Evidence from a Betting Exchange

Abstract

How does insider trading affect price efficiency? We show that insider trading substantially contributes to the creation of efficient markets by promoting quick price discovery and improving the accuracy of prices. We use live-tennis betting data to isolate the activity of traders who have earlier access to important information. We find that prices start updating before the information about the winner of a set becomes public, which indicates the existence of insiders. Specifically, we measure that roughly 60% of the full subsequent price reaction and 70% of the total increase in the price accuracy occur in the insider trading window.

JEL Classification: G14, L83

Keywords: Market efficiency, insider trading, event study, betting

1 Introduction

In theory, the effect of insider trading on the efficiency of security market prices is uncertain. One view proposes that insiders, i.e., traders with monopolistic access to any pricing-relevant information, contribute to the creation of efficient markets by quickly and accurately incorporating new information into prices (Manne, 1966; Carlton & Fischel, 1983). Another view proposes that the existence of insiders can cause outsiders not to trade: as outsiders will also stop collecting information, the market will be less informationally efficient (Fishman & Hagerty, 1992). Depending on which effect dominates, insider trading can have a positive or negative overall effect on the efficiency of security prices.

In the corporate world, insider trading is the practice of trading securities by corporate insiders, such as top officers and directors, who have privileged access to information about the true value of their firm's stock. Most financial regulators further distinguish between legal insider trading, which is based on non-material (public or nonpublic) information, and illegal insider trading, which is based on material nonpublic information. Although no statutory definition of "materiality" is available, material information is generally defined as information expected to significantly affect the stock price of a company, such as news about major acquisitions or changes in top management.

Despite a rich literature on price efficiency, direct empirical evidence on whether and to what extent insider trading based on material information affects the efficiency of security market prices is scarce due to the illegality of such behavior in most countries and the consequent lack of data. The vast majority of studies related to the efficiency of stock prices show that (1) self-reported (legal) insider trades correctly predict future stock performance, meaning that corporate insiders can gain abnormal profits (for a review, see Lakonishok & Lee, 2001) and that (2) stock prices react fully and swiftly after corporate news announcements.¹ However, these studies do not investigate material insider trades.

¹For example, see Ikenberry, Lakonishok, and Vermaelen (1995) for stock split announcements, Mitchell, Pulvino, and

Cornell and Sirri (1992), Meulbroek (1992), and Chakravarty and McConnell (1997) analyze the market’s reaction to illegal insider trades (their data originate from criminal and civil litigation reports) and show that insider trades lead to more rapid price discovery. However, Chakravarty and McConnell (1999) dispute these studies by showing, using a refined methodology to re-analyze the data from those previous studies, that the effect of insider trades on prices is not discernibly different from that of non-insider trades. Thus, there is generally little and mixed evidence on the topic.

In this article, we present an analysis of unique and naturally occurring field data that provide a rare opportunity to isolate the effect of insider trading on securities price efficiency: live (or in-play) tennis betting. In this context, a trader who can obtain information about the match earlier than others is an insider. Most traders follow the match on TV or via internet score feeds, and whenever new relevant information becomes observable, they update their bets online.² Our analysis exploits the inevitable technical delay in the transmission of match information from the stadium to end receivers: for example, the delay from filming to receiving a TV broadcast is *at least* five seconds (Kooij, Stokking, van Brandenburg, & de Boer, 2014; Hutchins, 2014; Brown & Yang, 2017). Due to this significant delay, the insider traders trading directly from the stadium, an activity that is officially forbidden, benefit from a fleeting informational advantage compared to other slower traders such as “live”-TV or “live”-scoreboard traders. Thus, we can attribute the price update in the first five seconds after important news events during the match to the activity of insiders.

Specifically, we analyze the price reaction after a player wins a *set*.³ Winning a set is an important step toward victory; therefore, observing in advance which player won the set constitutes a considerable informational advantage to insiders. We use high-frequency betting data from Betfair, one of the largest betting exchanges worldwide, on 141 men’s

Stafford (2004) for merger announcements, and MacKinlay (1997) for earnings announcements.

²We use trader as a synonym for a bettor, gambler, or wagerer.

³The Appendix A.1 provides a brief introduction to tennis scoring rules and terminology.

tennis singles matches played at two major professional tennis tournaments, the French Open and the Wimbledon Championships, over the 2009–2014 period.⁴ We complement this information with match data about players, courts, the winner of the match, match start and end times. Our dataset provides us with two practical advantages with respect to standard financial data. First and most important, our second-by-second data allow us to determine the event time (when new information is released) with high accuracy and thus to measure with precision the impact of insider trading activity on prices. Second, because the players take a short break after the end of each set (a low-information period), we do not have to address the problem of confounding events.

Using event study methods, we find that the cumulative abnormal returns averaged across the 365 event observations (*CAAR*) begin increasing immediately after the set events: the *CAAR* values at one second and five seconds following the event are 0.82% and 3.06%, respectively, and are significant at the 1% level. Most important, we show that the cumulative abnormal return during the first five insider trading seconds following the set events—when the TV viewers and other slower traders have not yet seen the event—accounts for roughly 60% of the full price reaction observed once the public receives the new information. We show that the impact of insider trading is even larger (roughly 80%) for unanticipated news events, such as *tie-break* events, when inside information is more valuable.

Furthermore, we find that the prediction accuracy of the betting prices on the match outcome increases after news events, suggesting that the market correctly interprets and incorporates the new information. Importantly, approximately 70% of this gain in prediction accuracy is attributed to the activity of insider traders. Finally, we estimate that a simple trading strategy implemented in the seconds after news events yields significant trading returns to insiders, ranging from 1% to 6% when applied to balanced matches.

In this article, we can uncover the impact of insider trading in a real-world environ-

⁴Table A.1 in the Appendix A.2 provides the full list of matches.

ment. To the best of our knowledge, this paper is the first to provide a clear setting for testing the question of how insider trading influences price efficiency. Overall, we provide important evidence that insider traders substantially contribute to higher price efficiency and thus market quality.

The remainder of this article is organized as follows. Section 2 reviews the literature. Section 3 describes our setting. Section 4 discusses the relationship between insider trading and price efficiency. Section 5 presents the data. Section 6 outlines our empirical methodology. Section 7 presents the results, and Section 8 concludes the article.

2 Literature Review

The debate over the advantages and disadvantages of insider trading is important, as it might influence the decision on whether and to what extent to regulate such trading in financial markets (Leland, 1992). Researchers and regulators have examined both the fairness and the economic implications of insider trading. Concerning the latter, which is the focus of our article, the key issue involves the assessment of the impact of insider trading on price efficiency. In this regard, two contradictory views have emerged.

The first view, pioneered by the work of Manne (1966) and extended by Carlton and Fischel (1983), claims that trading on inside information leads to more informationally efficient stock prices. The main argument of the supporters of insider trading is that insider trades are informative about the future performance of a firm and thus about its true value. Accurate stock prices help the capital market to efficiently allocate its resources: for example, a takeover decision is often based on a stock-price-based estimate of the target's value.

The second view claims that insider trading leads to less informationally efficient stock prices due to adverse selection costs (Fishman & Hagerty, 1992). According to this logic, some outsiders perceive the market to be unfair and thus might stop trading

and searching for relevant information because they anticipate that insiders still have superior knowledge. As a consequence, prices will be less informationally efficient and more volatile, and market liquidity will decrease (Leland, 1992).

According to Fama (1970), asset prices are efficient if they fully reflect all available information, public or private. Using event study methods, several financial studies have tested the semi-strong form of the efficient market hypothesis (EMH), which states that security prices swiftly and fully incorporate all relevant public information. The evidence is mostly supportive because stock prices adjust rapidly, usually within a day, to announcements (for an overview, see Fama, 1991).⁵ In general, event studies using stock returns face a number of issues, including the difficulty of determining the precise event date (the moment at which information is publicly disseminated), the effect of confounding events around the event date, and the need to use an asset pricing model to evaluate the normal stock returns.

Other studies test the strong form of the EMH, which states that prices include all private and public information, by investigating whether corporate insiders can gain abnormal profits when trading their company's stock (Lakonishok & Lee, 2001). Rozeff and Zaman (1988), for example, find evidence that managers profitably trade their company's shares before important corporate events and gain abnormal profits, suggesting that their trades are based on important information. However, these articles investigate *legal* insider trading, i.e., trading on *non*-material inside information that is self-reported by corporate insiders to the local financial authorities.⁶ A drawback of these studies lies in their inherent inability to precisely determine the motivation behind these trades: for example, managers might trade their firm's stock for hedging, risk-sharing, or liquidity reasons (Jaffe, 1974; Ke, Huddart, & Petroni, 2003).

⁵Some researchers have also observed a pre-announcement drift in stock prices (a so-called price run-up) and have suggested that this drift is caused by illegal insider trading on leaked material information (Mitchell et al., 2004).

⁶Corporate insiders must disclose their trades in their firms' securities to the local financial market authorities. In the U.S., the Securities and Exchange Commission (SEC) publishes monthly a list of all self-reported insider trades in the *Official Summary of Insider Transactions*.

Most relevant to price efficiency studies is illegal insider trading activity, i.e., trading on material inside information. However, corporate insiders obviously refrain from reporting violations of the law to authorities, and hence, the true amount of insider trading activity cannot be determined with precision (Keown & Pinkerton, 1981). Thus, despite the importance of the subject, the empirical evidence on the impacts of trading on material information on price efficiency is scarce.

One important exception is the study of Meulbroek (1992), who uses information on 183 illegal insider trading cases (mostly takeover-related events) detected by the Securities and Exchange Commission (SEC) to investigate the impact of these informed trades on the efficiency of stock prices. She finds that insider trading causes significant movements in prices, thereby increasing their accuracy. Notably, Meulbroek shows that the insider-induced price adjustment is almost 50% of the total price adjustment once the inside information is released to the public. Two further studies, Cornell and Sirri (1992) and Chakravarty and McConnell (1997), analyze illegal insider trades surrounding merger announcements at Campbell Taggart Inc. and Carnation Co., respectively. They show that, in these specific cases, insider trading led to more efficient stock prices.

However, Chakravarty and McConnell (1999) dispute these three previous studies on the grounds that they neglected to distinguish the effect of insider trading from that of non-insider trading on price discovery. Using a refined methodology, Chakravarty and McConnell re-analyze the data from those previous studies and show that the effect of insider trades on prices is not discernibly different from that of non-insider trades.

Some authors have used the data from betting markets to investigate the semi-strong form of the EMH for betting prices during sports events. For example, Croxson and Reade (2014) investigate the reaction of prices on a betting exchange to soccer goals scored within five minutes before the half-time break. Since there is little new information during half-time, they can test whether the prices quickly adjust to the news of the goal and remain constant thereafter, meaning that there is no price drift. The authors conclude that the

prices update swiftly and fully and that the betting market is semi-strong-form efficient. In contrast, Choi and Hui (2014) reject the semi-strong form of the EMH using very similar live soccer data because prices generally underreact to normal news but overreact to very surprising news.

Most related to our study is the work of Schnytzer and Shilony (1995), who also employ a setting in which they can reliably distinguish between the activity of two groups, one with and one without access to inside information. Schnytzer and Shilony use the horse pari-mutuel betting setting in Melbourne.⁷ Betting on horses is possible either on the racetrack using the pari-mutuel system (one pool) or the standard bookmaker system (composed of many different bookmakers offering fixed odds), or off the racetrack using the pari-mutuel system. In this setting, pari-mutuel traders at the racetrack have the advantage of observing how the bookmaker odds change shortly before the race. For example, they may observe that the bookmaker odds on a given horse have significantly decreased, suggesting that many traders, some of whom may have held inside information, favor that horse. Therefore, insiders profit from having a source of “second-hand” inside information via the bookmakers’ odds (Schnytzer & Shilony, 1995). In contrast, as communications to outside the stadium are prohibited by law, off-track pari-mutuel traders are not informed about last-minute developments and thus have an informational disadvantage. By analyzing the bets of these two isolated groups, Schnytzer and Shilony find that insiders make larger profits and that inside information on average better predicts the outcome of the race.

Finally, Brown (2012) analyzes the bid-ask spreads of the bookmakers’ odds during the 2009 Wimbledon men’s tennis final and asks whether some traders have an advantage due to superior analytical skills or due to access to material inside information. He observes an increase in the live bid-ask spreads prior to and during public information arrival and

⁷The pari-mutuel, or totalisator, is a betting system in which all the bets on the horses are combined in a so-called pool. The racetrack organizers handle the process by pooling and distributing the money, thereby earning a fee. The betting prices (or odds) on different horses vary until the pool is closed prior to the start of the event; at that moment, the payoff odds are determined.

attributes this to an increase in asymmetric information. Brown proposes that, in the presence of insiders, the bookmakers increase the bid-ask spread to offset losses to the informed traders. However, Brown provides no direct evidence on the effect of insider trading on price efficiency.

Overall, the empirical evidence on the economic effects of insider trading on market efficiency is limited and mixed in both the financial and betting markets. This article contributes to closing this literature gap by providing a simpler setting that allows us to overcome the difficulties of previous investigations.

3 Background Information

To understand the analyses in this paper, we introduce the advantages of using a betting exchange as a research laboratory and provide some important background information on live tennis betting. Most important, we discuss how we can distinguish between the trading activity of insider and outsider traders in our setting.

Compared to research in standard financial markets, which focuses primarily on stock prices, research in betting markets offers several advantages. First, whereas equity is infinitely lived, a betting asset (a bet) is short lived because the fundamental value of a bet is revealed at contract expiration. Therefore, one can test how accurately the prices of betting assets (the odds) predict an outcome, such as the winner of a match. Second, compared to stock prices, betting odds have two attractive properties: (1) new information has a larger impact on a bet's value because a bet has a high probability of default (Brown, 2012); (2) bet trading, especially during the match, is motivated primarily by information about the fundamental value of a bet (Vaughan Williams, 2009), whereas stock trading is motivated by a multitude of factors beyond the information about the fundamental value of equity, e.g., portfolio rebalancing or liquidity needs (Ke et al., 2003).⁸

⁸Uninformed or sentimental trading may occur in both markets.

Betting exchanges have several interesting features. In contrast to bookmakers, exchanges provide an online market for opinions where traders bet against other traders (without an intermediary) by offering and accepting odds under which they are willing to buy (*back*) or sell (*lay*) a certain bet. Thus, a betting exchange works very similarly to a traditional stock exchange. In this article, we use data from one of the major online betting exchanges worldwide, Betfair. The trading volumes on Betfair are very large: for example, 1.2 billion bets were placed on the exchange in 2014, corresponding to three million bets per day, resulting in a total matched volume of \$92 billion (Betfair, 2015). In the United Kingdom, the total revenue generated by tennis betting is second only to that generated by soccer betting (Townend, 2016).

Betfair revolutionized the sports betting industry in 2000 when it introduced the live (or in-play) feature, allowing traders to continuously place or accept bets as the competition unfolds. Tennis live betting is very popular for several reasons. First, trading in tennis is simpler than in other sports because there are just two contingencies for each point: a player wins or loses it. Traders mostly follow the match on TV, and whenever new information becomes observable, they update their bets or place new ones (Brown & Yang, 2017). Second, liquidity, which is indispensable for trading, is generally high. Third, points are scored frequently, and in the space of seconds, a match can swing from a match point for one player to a match point for the other player; this constant fluctuation in the odds facilitates trading (Landi, 2015). Finally, there are many betting opportunities because there are many tournaments throughout the year.

In the following two paragraphs, we briefly illustrate the basics of live betting. Fig. 1 depicts a screenshot of the live order book during the match between Murray and Wawrinka taken from *betfair.com* on 9 June 2017.⁹ The order book presents the bets available to bet on the match winner: these are the bets previously placed by other traders and that have not yet been matched. The third (fourth) column shows the current best

⁹This match, played at the French Open, serves only as illustration and is not included in our sample.

back (lay) odds for each player. The available volumes (in €) are provided below the odds. According to Fig. 1, the best odds to bet on Murray are 2.3, and it is possible to bet up to €13,256 at this price. As the inverse of the match odds can be interpreted as the player’s probability of winning the match, after four hours of play, Wawrinka was perceived as slight favorite for the victory, with a probability of roughly 57%.

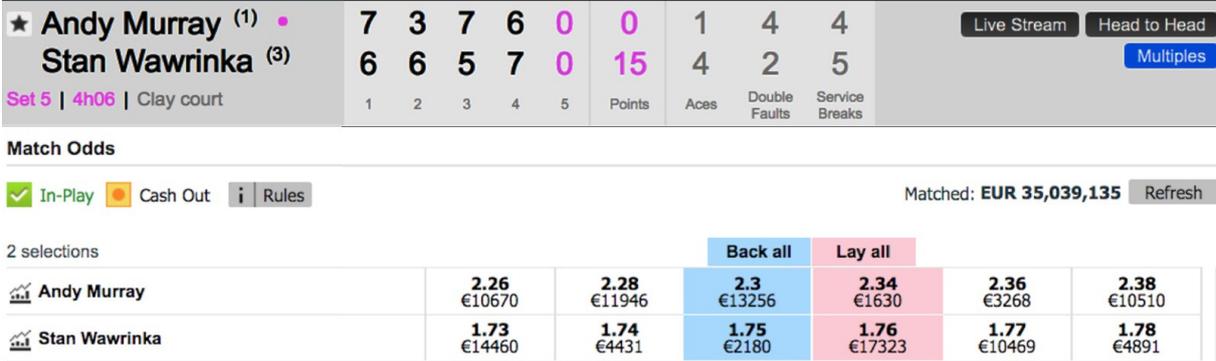


Figure 1
 Displayed is a screenshot from *betfair.com* illustrating its interface. The screenshot was taken on 9 June 2017 during the match between Murray and Wawrinka at the French Open tournament. Four hours into the game, the total matched amount was €35 million (approximately \$39 million).

The simplest way to bet on Murray is to click on 2.3 (under “Back all”) and enter the desired stake: by doing so, the trader (the backer) enters into a contract with one or more traders (the layers) who are taking the opposite position (against Murray) and offering odds of 2.3.¹⁰ The backer is placing a *market* order, thereby matching the outstanding orders previously placed by the layers. Assuming that the backer puts €10 at stake and Murray wins the match, the payoff to the backer is €13 plus the initial €10 stake.¹¹ As betting is a zero-sum game, the layer loses €13. If the backer wishes to back Murray at the higher odds of 2.38, the backer places a *limit* order. If entered, this order shows up on the “Lay all” side of the order book and can be matched by a trader willing to lay Murray at those odds.¹²

In tennis betting, a trader who can obtain information about the match earlier than

¹⁰Laying is an exclusive feature of betting exchanges in which layers take the role of a bookmaker offering the odds.
¹¹Here, we ignore the 5% commission on winning bets collected by Betfair (no commission is paid when a loss is incurred).
¹²The terms *market* and *limit* orders are also common in the standard financial markets. For example, placing a market order corresponds to either a buy or sell order at the current market price, whereas placing a limit order corresponds to either a buy order at a price below the current market price or a a sell order at a price above the current market price.

others is an insider. Match information about the score or the physical and mental status of the players is material because it will impact the prices. The fastest way for traders to access this information is to sit directly in the stadium, which also allows them to see special situations, such as signs of injury or the calls for physiotherapy assistance. As the venue where tennis is played is called the court, insider traders are sometimes referred to as courtsiders.¹³ Although trading from the stadium was generally legal under French and British law for the 2009–2014 period covered by our sample, such activity generally violates the terms and conditions of the ticket purchase and is therefore prohibited.¹⁴

Brad Hutchins, a former tennis courtsider, wrote the following in his book: “If you’re working on centre court at a Grand Slam [tennis tournament], there could be up to twenty other traders gambling on court. [...] When courts are competitive, speed becomes the key. [...] If everyone gambling at home is watching a television feed that is delayed by five seconds, you have just five seconds to take advantage. That’s plenty of time to log a point” (Hutchins, 2014, p. 28). Trading speed is essential for live traders, who have large incentives to quickly update their bets to gain a profit after each point. To trade algorithmically, some sophisticated courtsiders transmit live score data to remotely located computers, which then place new orders on the betting exchange (Dickson, 2015).

In contrast, off-stadium traders (outsiders) rely on communication technologies to stay informed about the progress of the match. The most comprehensive source of information for outsiders are TV broadcasts; they not only provide score information (as do internet scoreboards) but also key information such as a player’s fatigue level, confidence, or tactics. However, TV images are significantly delayed in comparison to the real match. In the terms and conditions for betting, Betfair clearly indicates that any transmissions described as “live” by some broadcasters may actually be delayed and that the extent

¹³Analogously, different terms exist for other sports, such as pitchsiding in cricket and soccer.

¹⁴For example, Article 24 of the ticket terms and conditions at Wimbledon (2017) reads: “Betting is prohibited in the Grounds at all times”, and Article 19 specifies that “[t]he use of photographic equipment, mobile telephones, computers, tablets or other electronic devices, communication devices, audio-visual equipment or radios must not [...] supply or transmit data for the purposes of betting or gambling (or assisting for these purposes).”

of any such delay may vary, depending on the set-up through which they are receiving pictures or data (Betfair, 2017).¹⁵ Due to the large number of factors determining the length of the TV delay, accurate estimates are difficult. Kooij et al. (2014) measure the delay in live TV broadcasts in the Netherlands. They estimate that a minimum delay of four seconds is introduced by the steps of encoding and modulation of the images, on top of which one should add between one and six seconds for the transmission.¹⁶

Because Kooij et al.'s minimal delay of five seconds corresponds with the TV delay mentioned by Brad Hutchins and by Betfair, and because other information sources such as internet scoreboards are also delayed (Landi, 2015), we use five seconds as the cutoff for distinguishing the insider trading period (within five seconds after the event) from the outsider trading period (from six seconds after the event onward). In other words, the communication delays allow us to differentiate insider from outsider trading activity: any price movement observed within five seconds following an event would reveal the presence of insider traders, if any.

In an effort to slow down insider traders, Betfair has implemented a bet processing delay, a so-called speed bump. The speed bump imposes a delay of five seconds between the time at which a market order (a new bet) is submitted and the time at which it is logged on the exchange and is eventually matched with outstanding orders (Brown & Yang, 2017). This mechanism is designed to protect the slower traders by giving them time to cancel or reduce the size of their outstanding (unmatched) orders without any

¹⁵On its webpage, Betfair writes the following: "Although the current score, time elapsed, video and other data provided on this site is sourced from "live" score feeds provided by third parties, you should be aware that this data may be subject to a time delay and/or be inaccurate. Please also be aware that other Betfair customers may have access to data that is faster and/or more accurate than the data shown on the Betfair site. If you rely on this data to place bets, you do so entirely at your own risk. Betfair provides this data AS IS with no warranty as to the accuracy, completeness or timeliness of such data and accepts no responsibility for any loss (direct or indirect) suffered by you as a result of your reliance on it."

¹⁶Concerning the transmission of TV images, the length of the delay can greatly vary, depending on the transmitting technologies (satellite is slower than terrestrial, while internet streaming is the slowest), the quality (high definition is slower than standard definition), the subscription (digital is generally slower than analog TV), and the broadcaster (which use different hardware). The delay increases with the geographical distance between the event, the local broadcaster, and the household (Kooij et al., 2014). Fig. 1 in Kooij et al. (2014) illustrates the various steps of a TV content delivery chain: live TV broadcasts differ from pre-recorded content (such as a TV series) because they introduce additional steps (and delays) due to the filming and transmission to the broadcaster from the event location. Furthermore, as the traders on a betting exchange are dispersed all over the world and use different technologies, an attempt to determine the actual average delay for all the traders is unrealistic.

delay.¹⁷ In practice, once the market trader has entered the stake and confirmed the order, Betfair begins a five-second countdown, at the end of which the market order is logged on the exchange; if the counterparty has not yet canceled the outstanding bet (within five seconds), the two orders are matched; if the counterparty has already canceled it, no transaction will take place.

However, the length of the speed bump may be inappropriate for protecting slower traders. First, as Kooij et al. (2014) shows, the actual TV delay may be well longer than five seconds. Moreover, the traders also need some time to react and update their bets. Second, Brown and Yang (2017) suggest that insider traders may circumvent the speed bump by placing two limit orders, one on each player, and then timely canceling the wrong position once the new score data come in. Since canceling a limit order occurs without delay, insider traders can still profit from their time advantage. Dan Dobson, a former insider trader, confirms this practice in an interview with the BBC: “We had an automated system whereby the point data would come in and then we would cancel any [outstanding] bets that we had in the market that we deemed were at the wrong price [...]. Then we would place bets straight back into the market that we deemed were now the correct price” (Cox, 2015). Clearly, such strategies (which we illustrate with an example below) reduce the protection offered by the speed bump to slower traders.

4 Insider Trading and Price Efficiency

Price efficiency studies investigate the degree to which asset prices timely and accurately reflect all available information. Our null hypothesis is that insider trading does not affect the efficiency of betting prices:

Hypothesis 1. Price efficiency is unaffected by insider trading.

¹⁷However, changing the odds of a limit order is treated as placing a new market order and is subjected to the speed bump.

In a first step, we analyze how quickly the market incorporates new information into the bets and what the role of the insiders is. Market efficiency theory (e.g., Manne, 1966; Carlton & Fischel, 1983) predicts that because insider trading is based on new, relevant information, insiders contribute to price discovery by rapidly incorporating information into market prices. Therefore, the alternative hypothesis is that insider trading is directly linked to the speed of incorporating new information into the prices of financial securities (Fernandes & Ferreira, 2009). We test our hypothesis using event study methods, which allow us to observe when and how the betting prices update following important informational events. Given the minimum communication delay of five seconds, we will reject our null hypothesis if we observe a significant price adjustment during the first five insider trading seconds following the event.

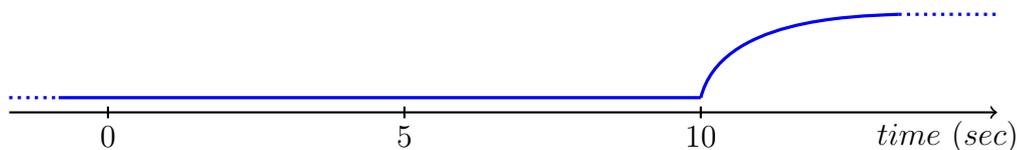
In the following, we lay out three broad scenarios regarding the hypothetical impact of insider trading on the price discovery. Fig. 2 depicts the earliest time at which the cumulative abnormal returns start increasing following an important event.¹⁸ The horizontal axis represents the time elapsed (in seconds) after the event, which takes place at time zero. In all scenarios, we assume a minimum informational delay (for the outsiders vs. the insiders) of five seconds.

Scenario 1 depicts a situation without insiders in which all traders are, for example, TV traders. The traders see the delayed TV images five seconds after the actual event ($time = 0$) and place their market bets. After a five-seconds delay due to the speed bump, the new orders are logged on the exchange. Thus, the price adjustment begins, at the earliest, ten seconds following the event.

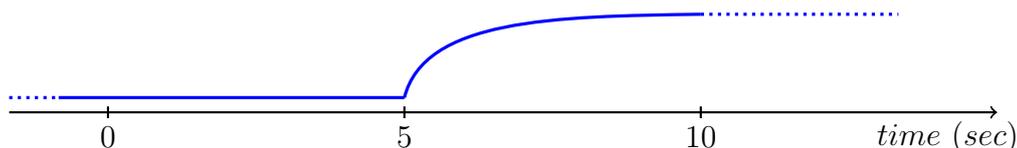
Scenario 2 depicts a situation with insiders and outsiders. If insiders observe the event as it happens ($time = 0$) and place new market orders that are subjected to the speed bump, the price adjustment begins, at the earliest, five seconds following the event.

¹⁸In Section 6, we describe the event study methodology in detail. The cumulative abnormal returns represent the sum of the abnormal returns (the difference between the realized returns and the expected normal returns) over a specified period after the event. To understand Fig. 2, one must understand that the price adjustment begins when the cumulative abnormal returns start to increase.

Scenario 1. Without insiders.



Scenario 2. With insiders.



Scenario 3. With insiders circumventing the speed bump.

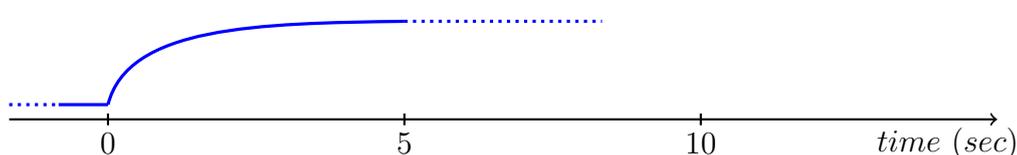


Figure 2

Displayed is the timeline of the event study under three scenarios. Time is in seconds, and zero corresponds to the second when a player wins a set. The line indicates the assumed time pattern of the cumulative average abnormal returns computed from the betting odds. The line is flat until new material information is released.

However, outsiders who see the event at second five may cancel, without any delay, their limit orders on the wrong side of the market. Therefore, under this scenario, we cannot distinguish the activity of insiders from that of outsiders.

Scenario 3 depicts the same situation as *Scenario 2*, with one exception: we assume that some sophisticated insider traders are able to circumvent the speed bump by timely canceling their limit orders on the wrong side of the market without any delay, as suggested by Brown and Yang (2017) and former courtsider Dan Dobson. For example, insiders observe Federer winning the set and immediately cancel their orders for “Federer to lose” but not those on “Federer to win”. Contemporaneously, the slower traders, who are still unaware of the outcome of the set, will see their market orders on “Federer to win” matched but not those on “Federer to lose”. Thus, under this scenario, the price adjustment begins immediately following the event. If the adjustment is significant, we will reject our null hypothesis.

Overall, we assume that insider trading has two effects on the timing of the price adjustment: (1) the presence of insiders advances the start of the adjustment by roughly five seconds (*Scenario 2*), which corresponds to the length of the informational delay; (2) the presence of some sophisticated insiders further advances the price adjustment by roughly five seconds, as those traders avoid the speed bump (*Scenario 3*).

In a second step, to investigate how accurately insiders incorporate new information into their bets, we analyze the change in the predictive ability of the odds during the first insider trading seconds. If insider trading increases the accuracy of the odds, the betting odds after the insider trades should better predict the final outcome, that is, the winner of the match. Regarding the three scenarios presented above, the prediction accuracy of the betting odds should start increasing after ten seconds under *Scenario 1*, after five seconds under *Scenario 2*, and immediately under *Scenario 3*. If the third scenario is observed and the insider trading contribution to the total change in price accuracy is substantial, we will reject *Hypothesis 1*.

5 Data

Our sample consists of 141 Grand Slam men's singles matches played between 2009 and 2014 at Roland Garros (61 matches) and at Wimbledon (80 matches). The betting data originate from Betfair and are provided by Fracsoft, the official data vendor. For every second of a match (in total, our sample includes 1,296,688 seconds), we have the best back and lay odds. The matched volume totals \$2.95 billion, 85% of which is generated live.

Detailed match data are provided by IBM, the official supplier of information technology to both tennis tournaments. Beyond general information about the match, such as players and the start and end match times, the IBM data also contain 31,018 point-level observations of the score, time (exact to the second), and winner. The score is fed into

the system directly by the match umpire using a computer, whereas other statistics are collected by analysts who attend the match and manually feed the data into the system.¹⁹

Table 1 provides summary statistics. Our final sample is heterogeneous, containing matches from the first stage up to the finals and a total of 83 unique players. The average live matched volume is roughly \$18 million. Usually, semifinal and final matches attract larger volumes than first-round matches, as shown by the \$70.5 million matched during the 2013 Wimbledon semifinal between Djokovic and Del Potro. Different from other sports, the speed bump in tennis has remained fixed over time at five seconds. An average match lasts 155 minutes and consists of 3.6 sets, 36 games, and 221 points.

Table 1
Summary statistics.

Variables	Mean	Std. dev.	Min.	Max.	<i>N</i>
Panel A: betting odds					
In-play matched volume (mio \$)	18.1	19.4	0.2	70.5	1,296,688
In-play order processing delay (seconds)	5	0	5	5	1,296,688
Panel B: match information					
Duration (minutes)	154.7	50.1	69.9	284.8	141
Number of points	221.2	66.1	117	437	141
Number of sets	3.6	0.7	3	5	141
Number of games (in match)	35.8	10.1	20	77	141

Notes: The table reports the summary statistics for the 141 men’s singles matches in our sample. The matches were played at the French Open and at the Wimbledon Championships between 2009 and 2014 (see Table A.1 in the Appendix). In Panel A, we report the statistics for the betting odds from Betfair. In Panel B, we report the match statistics from IBM.

We use the end of a set—when a player wins a set—as the informational event, as it satisfies three important criteria. First, the end of a set is an important news event because it is a pivotal moment in a match. At the French Open and at Wimbledon, the player who first wins three sets wins the match: therefore, a player’s probability of winning the match increases after winning a set. Second, the end of a set is an easily observable event. Third, a 120-second break follows the end of each set. This is advantageous because our event study results are less affected by overlapping and confounding events in the seconds following the end of a set.²⁰ Overall, our sample includes 365 set events. Since the betting

¹⁹The International Tennis Federation (ITF), the governing body of tennis, requires the umpires to “timely and accurately” enter the points in their computers throughout the match.

²⁰Crosson and Reade (2014) considers the 15-minute halftime break in soccer matches as a “low-information period” because important information is rarely revealed during the break. Similarly, in tennis, little information is revealed during the break, except when players start showing clear signs of fatigue, stress, or pain.

market is immediately closed when a player wins the match, we exclude the last set from our analyses.²¹

In a further analysis, we analyze a subsample encompassing the sets won after a tie-break ($N=79$). A tie-break decides the outcome of very balanced sets, and its outcome is often unanticipated. When we compare the outcome of a tie-break with the outcome of a set won by a large margin (e.g., six games to one), the tie-break usually has a larger surprise component. In our sample, a tie-break lasts on average approximately eight minutes (with a minimum of four minutes and a maximum of 17 minutes), and thus, the traders have enough time to recognize the importance of the moment and are prepared to adjust their bets.

Fig. 3 illustrates some characteristics of our two main events, the set and tie-break events, and provides insight into our data. Specifically, it shows the evolution of Del Potro's odds-implied winning probability (henceforth, WP) over the match when he played against, and lost to, Djokovic at Wimbledon on 5 July 2013.²² As the match started, Del Potro's WP was 15.3%. The vertical dashed lines indicate the end of a set (excluding the last one). Djokovic won the first set, resulting in a decrease in Del Potro's WP from 15.3% to approximately 8%. In the second part of the second set, Del Potro won a game when Djokovic was serving (a so-called break), resulting in a large increase in Del Potro's WP. Although Del Potro won the second set, his WP did not increase by much, perhaps because the market had already anticipated this outcome after the break.

The outcome of the third and fourth sets was highly uncertain and was decided only after a tie-break. Del Potro lost the third set, resulting in a large decrease in his WP from 33% (at the beginning of the tie-break) to approximately 12%. In the middle of the fourth set, Djokovic won a game when Del Potro was serving (a break), resulting in

²¹At the French Open and at Wimbledon, a match is played as the best of five sets. Since we do not include the last set when the match ends, we analyze between two and four set events per match. The "Number of sets" reported in Table 1 corresponds, however, to the statistics for the original sample before excluding the match points.

²²In a live betting market, the price throughout the match of a betting asset, such as "Del Potro to win the match", is represented by its live odds. By computing the inverse of these odds, we can derive the aggregate traders' belief about a given player's probability of winning the match (the WP) at any moment of the match (Hasbrouck, 1991).

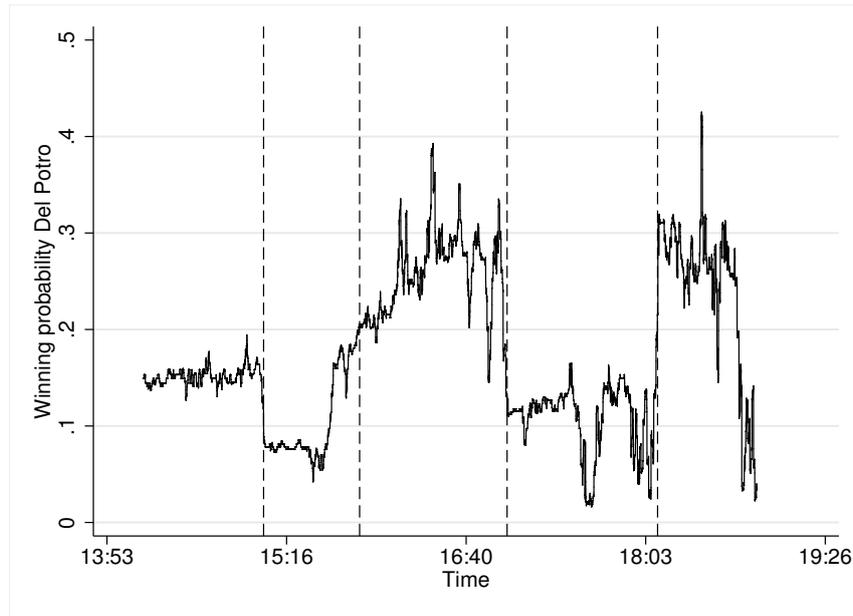


Figure 3

Displayed is Del Potro's odds-implied match winning probability for his 5 July 2013 semifinal match against Djokovic at Wimbledon. Del Potro lost in five sets after more than four hours of play with the following score: 5–7; 6–4; 7–6; 6–7; 3–6. A total of roughly \$74.8 million was bet on this match, 94% of which (\$70.5 million) was placed live. The vertical dashed lines indicate the end of a set (excluding the last one).

a large decrease in Del Potro's WP, which reached a record low of 3%. Del Potro won the fourth set by a small margin, resulting in a spike in his winning probability from 4% to approximately 30%. Overall, this example makes it apparent that the betting prices readily react to new information and that the price adjustments are larger after tie-breaks because their outcome is more uncertain.

6 Event Study Methods

The aim of our event study is to measure the betting odds reaction around the release of important information and to quantify the extent to which insider trading affects this reaction. Assuming rationality in the marketplace, security prices should timely and accurately reflect all available information (Fama, 1970).

We follow the classical methodology presented by MacKinlay (1997). Fig. 4 summarizes the timeline of the baseline event study. We define the moment when a player wins a

set or a tie-break as the event time ($\tau=0$). The event window spans from $\tau_2=-2$ to $\tau_3=14$ and lasts seventeen seconds ($L_2=17$). We account for the possibility that some traders anticipate the outcome of the point or that the umpire may update the score with a brief lag by letting our baseline event window begin two seconds before the event.

The estimation window spans from $\tau_0=-10$ to $\tau_1=-3$ and lasts eight seconds ($L_1=8$). This window does not begin immediately after the end of the previous point (roughly at $\tau=-20$) to avoid informational spillovers from the previous point. According to the ITF rules, the time between two consecutive points should be 20 seconds—in our sample, the median elapsed time is 22 seconds. Furthermore, as the event and estimation window should not overlap, the estimation window ends at $\tau_1=-3$.²³

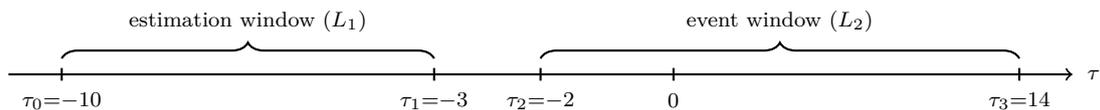


Figure 4

Displayed is the timeline of the baseline event study, where τ represents the time in seconds during a match. The news event takes place at $\tau=0$. The estimation window length is $L_1=8$, and the event window length is $L_2=17$.

Because we want to analyze the reaction of the odds on the player who won the set, we consider only the odds of the set and tie-break winner. For each second τ surrounding an event, we compute the average mid-odds from the best back ($back_\tau$) and sell odds (lay_τ) as follows:

$$mid_\tau = \frac{back_\tau + lay_\tau}{2} . \quad (1)$$

The mid-odds, which are expressed in decimal terms, can range from a minimum of 1.01 to a maximum of 1000. Then, we derive the implied winning probability by taking the inverse of the the mid-odds ($WP_\tau = \frac{1}{mid_\tau}$), which approaches 99% when the mid-odds approach 1.01. Finally, we compute the actual returns from the odds-implied winning

²³In a robustness check, we also test an alternative estimation window lasting 15 seconds, from $\tau_0=-17$ to $\tau_1=-3$ jointly with an alternative event window lasting 19 seconds, from $\tau_2=-2$ to $\tau_3=17$. The results, which we present in Table A.3 in the Appendix A.3, are similar.

probability as follows:

$$R_\tau = \frac{WP_\tau}{WP_{\tau-1}} - 1 . \quad (2)$$

When a player wins the set or tie-break, the odds should adjust downward, causing the odds-implied winning probability to increase and resulting in a positive return.

Since no asset valuation model is available to measure abnormal returns in tennis, we use the mean return over the estimation window as an estimate for the normal return. We estimate the normal return by calculating the arithmetic mean of the returns over the estimation window for each event j in our sample as follows:

$$\bar{R}_{j\tau} = \frac{1}{L_1} \sum_{\tau=\tau_0}^{\tau_1} R_{j\tau} , \quad (3)$$

where L_1 is the estimation window length. We then compute the abnormal returns for each event j by subtracting the normal returns from the actual returns over the event window: $AR_{j\tau} = R_{j\tau} - \bar{R}_{j\tau}$. Under the null hypothesis, the abnormal returns will be jointly normally distributed with a zero conditional mean and conditional variance $\sigma^2(AR_j)$. It can be assumed that $\sigma_{AR_j}^2 = \sigma_{\epsilon_j}^2$, where $\sigma_{\epsilon_j}^2$ is estimated by computing the sample variance of the returns over the estimation window as in (MacKinlay, 1997).

Finally, we aggregate the abnormal returns along two dimensions. First, we aggregate the returns along the event dimension to generate the average abnormal returns AAR_τ . Second, we aggregate the average abnormal returns along the time dimension to compute the cumulative average abnormal returns $CAAR_{(\tau_2, \tau_3)} = \sum_{\tau=\tau_2}^{\tau_3} AAR_\tau$.

7 Empirical Results

7.1 Event Study

Fig. 5 presents the cumulative average abnormal returns ($CAAR$) and the corresponding 95% confidence interval over the event window for the 365 set events.²⁴ We observe a significant positive reaction in the betting odds in the first seconds immediately following the set event. Most important, during the first five insider trading seconds, the $CAAR$ is 3.06% and is statistically significant at the 1% level. Therefore, some sophisticated insider traders are able to circumvent the speed bump by timely canceling, without any delay, their limit orders on the wrong side of the market, as described in *Scenario 3*.

The pattern displayed in Fig. 5 shows that new information is fully incorporated into the betting odds within six to seven seconds, after which the $CAAR$ stabilizes around 5%. Thus, insider trading accounts for more than 60% of the full price reaction observed once the public (the outsiders) receives the information about the set winner. This result shows that insider trading substantially contributes to advance the beginning of the price discovery process, which contrasts with *Hypothesis 1*.

Furthermore, Fig. 5 shows an average abnormal return (AAR) of 1.44%, the largest over the event window, at second six. Apparently, insider traders not only cancel the wrong outstanding limit orders but also immediately place new market orders (which undergo the five-seconds speed bump) at the price that they deem correct to increase their profits. By slowing down insider traders, the speed bump somewhat lengthens the price discovery process by five seconds.

Fig. A.1 (in the Appendix A.3) presents the $CAAR$ and the corresponding 95% confidence interval over the event window for the 79 tie-break events (i.e., sets won at the tie-break).²⁵ Overall, the $CAAR$ pattern is similar to that displayed in Fig. 5, with three main differences: first, the total adjustment of approximately 8.8% is significantly

²⁴In the Appendix A.3, Panel A in Table A.2 reports the full results.

²⁵In the Appendix A.3, Panel B in Table A.2 reports the full results.

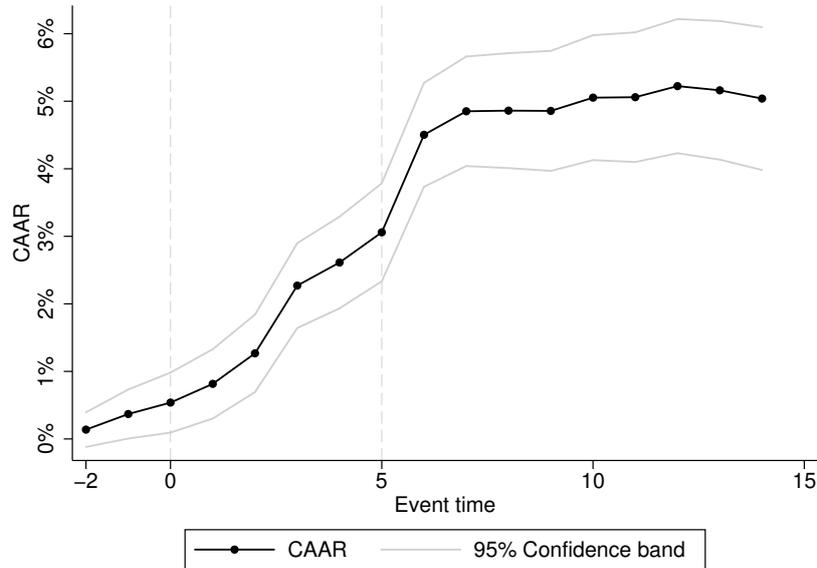


Figure 5
Cumulative average abnormal returns (*CAAR*) for the set events.

larger than for the set events, which was roughly 5%. This is due to the importance of winning a tie-break in a balanced match and the difficulty of predicting its outcome. Second, the largest average abnormal return (*AAR*) of 3.18% now takes place at $\tau=3$. This supports Hutchins’s (2014) assertion that, during pivotal moments, insiders attempt to trade even faster to gain an edge over other insiders. Third, insider trading during the first five seconds causes approximately 80% of the full price reaction, which is larger than the corresponding figure for the set events.

Taken together, our results show that the speed and amount of price discovery are affected by insider trading, which contradicts *Hypothesis 1*. Furthermore, the tie-break results show that the impact of insider traders on price discovery is even larger during important moments, when inside information is more valuable.

The drawback of any event study methodology is that the economic interpretation of the results depends on the underlying assumptions used to estimate normal returns and for statistical testing. The null hypothesis that the average abnormal returns are zero implicitly includes a test of whether the model used for measuring normal returns is

correct—this issue is known as the joint hypothesis problem (Fama, 1991). For example, stock efficiency studies rely on some equilibrium asset pricing models to measure any normal return, i.e., the expected return given the absence of a particular event such as an earnings announcement. Thus, any abnormal return observed may reflect market inefficiency, an inaccurate pricing model, or both (Fama, 1991). To ensure that our results are not affected by our estimation of the normal returns, we repeat the event study by assuming that the betting prices would not change and that the normal returns would be zero over the event window for all set events. The (untabulated) results of this robustness check are qualitatively the same.

7.2 Prediction Accuracy Analysis

In this subsection, we examine how insider trading affects the level of predictive ability and thus the efficiency of the betting odds. In contrast to most standard financial assets (e.g., stocks), the fundamental value of each bet is unambiguously revealed when the outcome of the underlying event—in our case, the winner of the match—is observed (Brown & Yang, 2017). The actual outcome serves as an indisputable benchmark against which the prediction accuracy of the odds can be tested (Vaughan Williams, 2009). Since a player either wins or loses the match, the difference between the actual outcome and his odds-implied winning probability at any second during the match can be interpreted as a “prediction error” (Camerer, 1989). As an illustration, if, after the second set, the market predicts Federer to win with a probability of 90% and he actually wins the match, the prediction error at that moment is 10%.

Following Wolfers and Zitzewitz (2004), we compute the absolute prediction error $APE_{j,\tau}$ for each event j in our sample over each event window as follows:

$$APE_{j,\tau} = |outcome_j - WP_{j,\tau}| , \tag{4}$$

where $outcome_j$ is a dummy variable, which takes a value of one for all the events where the winner of the set is the winner of the match and zero otherwise. The odds-implied winning probability of the player who wins the set is denoted by WP_τ and computed by taking the inverse of the mid-odds on the set winner ($WP_\tau = \frac{1}{mid_\tau}$). The APE ranges from zero for the most accurate prediction to one for the least accurate prediction.

According to *Hypothesis 1*, if insider trading has no effect on the price efficiency of the betting odds, the APE should not change significantly during the first insider trading seconds. In contrast, if insider trading has an effect on efficiency, we should observe a significant decrease in the mean APE in the insider trading period, indicating an increase in the accuracy of the odds in predicting the match outcome.²⁶

Thus, we plot the mean APE over the event window to examine the change in prediction accuracy after the set events. Fig. 6 shows that the mean APE decreases in the ten seconds following the set events, suggesting that the market correctly interprets the new information. Most important, we measure that roughly 70% of the total downward adjustment of the mean APE takes place in the first five insider trading seconds.²⁷ As the mean APE decreases, the prediction accuracy of the betting odds after the event increases. This suggests that insider traders on average correctly integrate the new information into their bets, thereby making the odds more efficient.²⁸ Thus, we reject *Hypothesis 1*.

²⁶For example, we consider a match won by Player 1. If Player 1 wins the first set, the WP should increase after the set event, as winning a set is an important step toward victory. Since the WP shifts closer to the fundamental value of one (*outcome*: Player 1 wins the match), the APE will decrease. In contrast, if Player 2 wins the first set, the WP should increase after the set event. Since WP shifts further away from the true fundamental value of zero (*outcome*: Player 2 loses the match), the APE will increase. However, the mean APE over all events is expected to decrease if insider trades on average correctly price the new information.

²⁷Concerning the tie-break events, Fig. A.2 in the Appendix A.3 shows that the mean APE decreases by 44.3%.

²⁸Other measures employed in the prediction accuracy literature such as the Brier score (see e.g., Franck, Verbeek, & Nüesch, 2010) or the Root Mean Square Error (RMSE) (see e.g., Spann & Skiera, 2003) lead to qualitatively similar results.

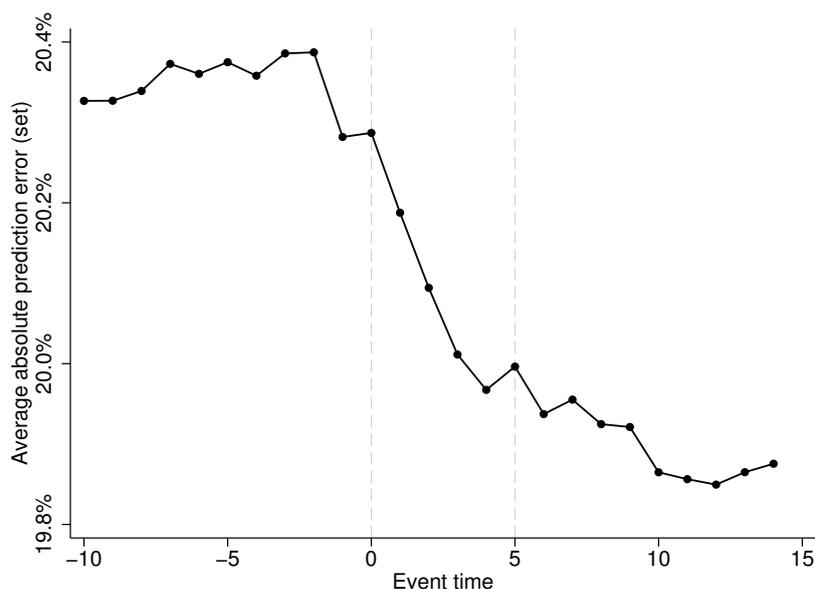


Figure 6
Mean absolute prediction error (*APE*) for the set events.

7.3 Trading Strategy

In a final step, we analyze whether insider traders can earn positive returns by betting from the stadium. By knowing before others in which direction the betting odds are likely to move, an insider may use a simple “sequential back-lay” trading strategy. Accordingly, the insider first backs the player who wins a set and, a few seconds later, lays that same player. By taking both sides of the market (sequentially betting for and against the same player), the insider trader will earn a certain payoff regardless of the final outcome of the match.²⁹ The payoff will be positive if the back-lay odds difference is positive, and larger if that difference is larger.

For example, an insider stakes \$100 to back the player (P1) who wins the set at odds of 2.0; after ten seconds, he or she stakes \$105 to lay P1 at the lower odds of 1.90. The gross return that this strategy locks in is 5%, regardless of the final outcome: if P1 wins the match, the back-bet win is \$100 and the lay-bet loss is \$95, yielding a gross profit of

²⁹Hedging a trade by taking the opposite position on the same player is a strategy called “greening up”. The name derives from the fact that, by dynamically backing and laying the same player correctly, the trader makes a positive profit (indicated by a green number on the trading software) regardless of the outcome.

\$5 (or a net profit of \$4.75 after deducting the 5% commission on winning bets); if P1 loses the match, the back-bet loss is \$100 and the lay-bet win is \$105, yielding exactly the same profit.³⁰ In this example, the market rapidly adjusted to the new information, and the odds for the set winner decreased (from 2.0 to 1.9) to reflect the increase (from 50% to 52.6%) in P1's probability of winning the match.

To assess the return that an insider trader could have earned by using the balanced sequential back–lay trading strategy, we initially assume that the insider backs the player who wins a set (at $\tau=0$) and lays the same player after exactly ten seconds (at $\tau=10$).³¹ Such a simple, automated trading rule yields -0.87% (p -value=0.09) for the 365 set events and -0.65% (p -value=0.69) for the 79 tie-break events. The returns are on average negative probably for two reasons: first, when one player is a strong favorite, the odds do not move by much (e.g., from 1.15 to 1.14) or at all after a set event because the traders anticipate its outcome; second, when an underdog wins the set, we compute the returns using the back/lay prices on him, but, due to very low volumes, the odds tend to be erratic.³²

As Landi (2015) suggests, such a trading strategy should yield higher returns when applied to balanced matches because the odds fluctuate more after the set events. Thus, we calculate the returns for the same strategy but limiting it to set events or tie-break events where the match is balanced, i.e., the odds-implied probability of the set winner is between 25% and 75% or, alternatively, between 40% and 60%. Since a trader can observe at any time if the match is balanced, this strategy is applicable in practice. After applying the “25-75” filter to the 365 set and 79 tie-break events, we find that a balanced back-lay strategy yields 1.48% (p -value=0.15) for 97 set events and 4.03% (p -value=0.03) for 29 tie-break events. Using the stricter “40-60” filter, the same strategy yields 2.96% (p -value=0.02) for 44 set events and 6.6% (p -value=0.11) for 12 tie-break events. As

³⁰The precise lay dollar amount yielding a so-called balanced trade is computed by multiplying the back amount by the back odds and dividing by the lay odds.

³¹We choose a window of ten seconds because Fig. 5 shows that after on average ten seconds, the odds stabilize. Theoretically, if the trader expects the lay odds to further decrease after ten seconds, he or she could wait longer before laying, resulting in a larger payoff. Here, for simplicity, we ignore the speed bump and assume that the new market orders are immediately matched.

³²This happens in particular in first-round matches.

expected, the trading profits are larger after tie-break events because these matches are usually more balanced and the liquidity is higher.

Overall, these results show that insiders can make substantial profits, especially during balanced matches and after tie-break set events, when information is more valuable. More refined trading strategies may yield even higher returns if applied at the right moment. However, one should not forget that, in practice, insiders have to cover their travel, stay, and entrance costs. Furthermore, other factors such as the competition with other insider traders, the presence of stadium security agents, low liquidity, and own trading mistakes may shrink the profits from such activity.

8 Concluding Remarks

In most financial markets, trading on material information by corporate insiders is against the law. Thus, the evidence on the extent to which this important form of trading improves price efficiency is, at best, scarce. We attempt to fill this gap in the literature by using a previously unexplored setting: live tennis betting. Due to inevitable communication delays, betting traders in the stadium have an informational advantage because they can observe material information about the match before other traders, such as TV traders. Thus, this setting offers a unique opportunity to distinguish the live trading activity of insiders from that of outsiders.

Our baseline event study uses 365 set events from 141 men's singles matches at two major professional tennis tournaments and shows that the betting prices rapidly adjust in the first five insider trading seconds following important moments—the end of sets and tie-breaks—when outsiders have not yet received the information. This result is surprising because courtsiding is officially prohibited by tournament organizers. Most important, we show that the price discovery is substantially affected by insider trading: the cumulative abnormal return during the insider trading seconds is between 60% and 80% of the full

price reaction observed once the public receives the new information. This result is larger than that of Meulbroek (1992), who finds, using illegal insider trading SEC data, that the insider-induced price adjustment accounts for almost 50% of the total price adjustment once the inside information is officially released. We also show that the the impact of insider trading is even larger for unanticipated news events, such as tie-break set events, when inside information is more valuable.

Furthermore, we find that on average insider trades help to improve the prediction accuracy of the prices. We find that the mean prediction error of the prices decreases after the news events, suggesting that the market correctly interprets the new information. Importantly, roughly 70% of this gain in prediction accuracy is attributed to the activity of insider traders. Finally, we estimate that a simple trading strategy implemented in the first seconds after the events yields significant trading returns to insiders, ranging roughly from 1% to 6% for balanced matches.

Our results are important, as they provide empirical evidence on the advantages of insider trading in terms of its contribution to price efficiency. Overall, our findings are supportive of the view that insider trading not only promotes quick price discovery, but it also improves the accuracy of the betting prices. These results support Manne's (1966) assertion that insider trading fosters efficient capital markets.

As betting on a betting exchange is a zero-sum game, one could also discuss how the insider activity disadvantages the slower traders. Although Betfair has implemented a speed bump in an effort to protect these outsiders from the insiders, our data suggest that sophisticated insiders have developed strategies to circumvent it. Further research should determine whether slower traders adopt more passive trading strategies or even stop trading when they anticipate that they cannot compete with insiders on speed. Overall, when debating the effects of insider trading on a betting exchange or on any financial market, these adverse selection costs to slower traders should be weighed against the positive externalities of greater price efficiency.

A Appendix

A.1 Tennis

This appendix introduces the basic rules of tennis and its jargon. These rules are specific to Grand Slam tournaments and can be found on the website of the International Tennis Federation. At the beginning of the match, a coin toss decides which player starts serving in the first game. Player 1 begins the match by serving in the first game of the first set—player 2 is the receiver. A player wins a point, sometimes referred to as point game, if the opponent cannot return the ball. A game is won when one player wins four points with a two-point difference, or when there is a two-point difference after a deuce, i.e., a score of 40–40 (3 points to 3 points in a game). A player has a gamepoint if he needs one more point to win the game: if this player is the receiver, the situation is called a breakpoint. A break (of service) happens when the receiver wins the game.

The players alternate serving every game, and they change ends after every odd-numbered game. A *set* is won when a player either wins six games with a two-game difference, or, in the case of a tie-break, when the score for one player is 7:6. A player has a setpoint if he needs one more point to win a set: if this player is the receiver, he has a breakpoint to win the set.

The *tie-break* begins when the game score is tied at 6:6 and is played until one player wins seven points with a two-point difference, or until there is a two-point difference when the point score is 6:6. At Grand Slam tournaments, a tennis match is played as the best of five sets, meaning that the first player winning three sets wins the match. The fifth set does not have a tie-break; the set (and match) is won when one player wins two more games than the other player.

A.2 List of the Matches

Table A.1
List of the matches included in our sample.

Player 1	Player 2	Date	Stage	Player 1	Player 2	Date	Stage
French Open:				Wimbledon (cont.):			
R. Gasquet	R. Stepanek	23-May-2011	1	A. Roddick	A. Murray	3-Jul-2009	6
A. Clement	M. Berrer	26-May-2011	2	A. Roddick	R. Federer	5-Jul-2009	7
J. Tipsarevic	R. Federer	27-May-2011	3	R. Federer	A. Falla	21-Jun-2010	1
S. Darcis	G. Monfils	27-May-2011	3	R. Federer	I. Bozoljac	23-Jun-2010	2
R. Gasquet	N. Djokovic	29-May-2011	4	R. Federer	A. Clement	25-Jun-2010	3
G. Simon	R. Soderling	30-May-2011	4	P. Petzschner	R. Nadal	26-Jun-2010	3
G. Monfils	R. Federer	31-May-2011	5	N. Djokovic	L. Hewitt	28-Jun-2010	4
A. Murray	J. Chela	1-Jun-2011	5	R. Federer	J. Melzer	28-Jun-2010	4
R. Nadal	R. Soderling	1-Jun-2011	5	R. Federer	T. Berdych	30-Jun-2010	5
R. Federer	N. Djokovic	3-Jun-2011	6	R. Soderling	R. Nadal	30-Jun-2010	5
R. Nadal	A. Murray	3-Jun-2011	6	J. Tsonga	A. Murray	30-Jun-2010	5
R. Nadal	R. Federer	5-Jun-2011	7	A. Murray	R. Nadal	2-Jul-2010	6
M. Llodra	G. Garcia-Lopez	28-May-2012	1	T. Berdych	N. Djokovic	2-Jul-2010	6
I. Sjsjing	G. Muller	27-May-2012	1	T. Berdych	R. Nadal	4-Jul-2010	7
M. Berrer	J. Melzer	27-May-2012	1	M. Kukushkin	R. Federer	21-Jun-2011	1
R. Federer	A. Ungur	30-May-2012	2	A. Mannarino	R. Federer	23-Jun-2011	2
J. Del Potro	M. Cilic	1-Jun-2012	3	M. Baghdatis	N. Djokovic	25-Jun-2011	3
F. Fognini	Jw. Tsonga	1-Jun-2012	3	D. Nalbandian	R. Federer	25-Jun-2011	3
G. Simon	S. Wawrinka	1-Jun-2012	3	A. Murray	R. Gasquet	27-Jun-2011	4
J. Monaco	R. Nadal	4-Jun-2012	4	M. Llodra	N. Djokovic	27-Jun-2011	4
J. Tipsarevic	N. Almagro	4-Jun-2012	4	M. Youzhny	R. Federer	27-Jun-2011	4
R. Federer	D. Goffin	3-Jun-2012	4	R. Nadal	M. Fish	29-Jun-2011	5
N. Djokovic	A. Seppi	3-Jun-2012	4	J. Tsonga	R. Federer	29-Jun-2011	5
N. Almagro	R. Nadal	6-Jun-2012	5	R. Nadal	A. Murray	1-Jul-2011	6
D. Ferrer	A. Murray	6-Jun-2012	5	J. Tsonga	N. Djokovic	1-Jul-2011	6
R. Federer	J. Del Potro	5-Jun-2012	5	R. Nadal	N. Djokovic	3-Jul-2011	7
N. Djokovic	J. Tsonga	5-Jun-2012	5	R. Federer	A. Ramos	25-Jun-2012	1
D. Ferrer	R. Nadal	8-Jun-2012	6	R. Federer	F. Fognini	27-Jun-2012	2
N. Djokovic	R. Federer	8-Jun-2012	6	R. Federer	J. Benneteau	29-Jun-2012	3
M. Raonic	M. Llodra	29-May-2013	2	R. Federer	X. Malisse	2-Jul-2012	4
G. Monfils	E. Gulbis	29-May-2013	2	R. Federer	M. Youzhny	4-Jul-2012	5
M. Przysieszny	R. Gasquet	31-May-2013	2	D. Ferrer	A. Murray	4-Jul-2012	5
N. Djokovic	G. Pella	30-May-2013	2	A. Murray	J. Tsonga	6-Jul-2012	6
V. Troicki	M. Cilic	31-May-2013	3	N. Djokovic	R. Federer	6-Jul-2012	6
N. Davydenko	R. Gasquet	1-Jun-2013	3	R. Federer	A. Murray	8-Jul-2012	7
T. Haas	J. Isner	1-Jun-2013	3	V. Hanescu	R. Federer	24-Jun-2013	1
K. Anderson	D. Ferrer	2-Jun-2013	4	R. Nadal	S. Darcis	24-Jun-2013	1
R. Nadal	K. Nishikori	3-Jun-2013	4	N. Djokovic	F. Mayer	25-Jun-2013	1
J. Tsonga	R. Federer	4-Jun-2013	5	D. Ferrer	M. Alund	25-Jun-2013	1
T. Robredo	D. Ferrer	4-Jun-2013	5	A. Ramos	J. Del Potro	25-Jun-2013	1
R. Nadal	S. Wawrinka	5-Jun-2013	5	J. Tsonga	E. Gulbis	26-Jun-2013	2
N. Djokovic	T. Haas	5-Jun-2013	5	F. Verdasco	J. Benneteau	26-Jun-2013	2
D. Ferrer	Jw. Tsonga	7-Jun-2013	6	N. Djokovic	B. Reynolds	27-Jun-2013	2
N. Djokovic	R. Nadal	7-Jun-2013	6	R. Gasquet	G. Soeda	27-Jun-2013	2
R. Nadal	D. Ferrer	9-Jun-2013	7	J. Levine	J. Del Potro	27-Jun-2013	2
V. Estrella Burgos	J. Janowicz	25-May-2014	1	J. Melzer	S. Stakhovskiy	28-Jun-2013	3
A. Dolgopolov	A. Ramos	25-May-2014	1	J. Janowicz	J. Melzer	1-Jul-2013	4
G. Elias	D. Schwartzman	25-May-2014	1	M. Youzhny	A. Murray	1-Jul-2013	4
R. Gasquet	B. Tomic	27-May-2014	1	N. Djokovic	T. Berdych	3-Jul-2013	5
G. Monfils	V. Hanescu	27-May-2014	1	F. Verdasco	A. Murray	3-Jul-2013	5
I. Karlovic	K. Anderson	31-May-2014	3	N. Djokovic	J. Del Potro	5-Jul-2013	6
E. Gulbis	R. Federer	1-Jun-2014	4	J. Janowicz	A. Murray	5-Jul-2013	6
T. Berdych	J. Isner	1-Jun-2014	4	N. Djokovic	A. Murray	7-Jul-2013	7
G. Garcia-Lopez	G. Monfils	2-Jun-2014	4	G. Dimitrov	R. Harrison	23-Jun-2014	1
M. Raonic	N. Djokovic	3-Jun-2014	5	A. Murray	D. Goffin	23-Jun-2014	1
T. Berdych	E. Gulbis	3-Jun-2014	5	P. Lorenzi	R. Federer	24-Jun-2014	1
G. Monfils	A. Murray	4-Jun-2014	5	S. Wawrinka	J. Sousa	24-Jun-2014	1
R. Nadal	D. Ferrer	4-Jun-2014	5	G. Dimitrov	L. Saville	25-Jun-2014	2
E. Gulbis	N. Djokovic	6-Jun-2014	6	A. Murray	B. Rola	25-Jun-2014	2
R. Nadal	A. Murray	6-Jun-2014	6	G. Muller	R. Federer	26-Jun-2014	2
R. Nadal	N. Djokovic	8-Jun-2014	7	L. Rosol	R. Nadal	26-Jun-2014	2
Wimbledon:							
G. Garcia-Lopez	R. Federer	24-Jun-2009	2	S. Wawrinka	Y-H. Lu	26-Jun-2014	2
P. Kohlschreiber	R. Federer	26-Jun-2009	3	N. Djokovic	G. Simon	27-Jun-2014	3
S. Wawrinka	J. Levine	27-Jun-2009	3	S. Giraldo	R. Federer	28-Jun-2014	3
A. Murray	S. Wawrinka	29-Jun-2009	4	M. Kukushkin	R. Nadal	28-Jun-2014	3
R. Soderling	R. Federer	29-Jun-2009	4	S. Wawrinka	D. Istomin	30-Jun-2014	3
T. Haas	N. Djokovic	1-Jul-2009	5	A. Murray	G. Dimitrov	2-Jul-2014	5
L. Hewitt	A. Roddick	1-Jul-2009	5	S. Wawrinka	R. Federer	2-Jul-2014	5
I. Karlovic	R. Federer	1-Jul-2009	5	N. Djokovic	G. Dimitrov	4-Jul-2014	6
T. Haas	R. Federer	3-Jul-2009	6	R. Federer	M. Raonic	4-Jul-2014	6
				N. Djokovic	R. Federer	6-Jul-2014	7

Notes: The table lists all 141 matches played at the French Open (Roland Garros) and at the Wimbledon Championships included in our sample. Stage indicates the tournament stage, from a first-round match (Stage=1) up to the final match (Stage=7).

A.3 Further Results

Table A.2
Event study results (baseline).

Panel A: set events ($N=365$)						
Time (τ)	<i>AAR</i>	t-statistic (θ_1)	p-value	<i>CAAR</i>	t-statistic (θ_2)	p-value
-2	0.14%	1.05	0.293	0.14%	1.05	0.293
-1	0.23%	1.76	0.077	0.37%	1.99	0.046
0	0.17%	1.29	0.194	0.54%	2.37	0.017
1	0.28%	2.11	0.034	0.82%	3.11	0.001
2	0.45%	3.45	0.000	1.27%	4.33	0.000
3	1.00%	7.66	0.000	2.27%	7.08	0.000
4	0.34%	2.61	0.008	2.61%	7.54	0.000
5	0.45%	3.41	0.000	3.06%	8.26	0.000
6	1.44%	11.03	0.000	4.50%	11.46	0.000
7	0.35%	2.66	0.007	4.85%	11.72	0.000
8	0.01%	0.07	0.941	4.86%	11.19	0.000
9	0.00%	-0.02	1.023	4.86%	10.71	0.000
10	0.20%	1.49	0.134	5.05%	10.7	0.000
11	0.01%	0.05	0.958	5.06%	10.33	0.000
12	0.16%	1.25	0.209	5.22%	10.3	0.000
13	-0.06%	-0.47	1.365	5.16%	9.86	0.000
14	-0.12%	-0.93	1.649	5.04%	9.33	0.000

Panel B: tie-break events ($N=79$)						
Time (τ)	<i>AAR</i>	t-statistic (θ_1)	p-value	<i>CAAR</i>	t-statistic (θ_2)	p-value
-2	-0.13%	-0.37	1.289	-0.13%	-0.37	1.289
-1	0.63%	0.97	0.082	0.49%	0.97	0.334
0	0.43%	1.47	0.24	0.92%	1.47	0.143
1	0.93%	2.55	0.01	1.85%	2.55	0.011
2	0.80%	3.26	0.028	2.65%	3.26	0.001
3	3.18%	6.56	0.000	5.82%	6.56	0.000
4	0.66%	6.76	0.068	6.48%	6.76	0.000
5	0.70%	7.00	0.054	7.18%	7.00	0.000
6	0.63%	7.18	0.081	7.81%	7.18	0.000
7	1.24%	7.89	0.000	9.05%	7.89	0.000
8	-0.25%	7.32	1.509	8.80%	7.32	0.000
9	-0.27%	6.79	1.537	8.53%	6.79	0.000
10	0.89%	7.20	0.014	9.42%	7.20	0.000
11	-0.27%	6.74	1.538	9.15%	6.74	0.000
12	-0.26%	6.33	1.526	8.89%	6.33	0.000
13	-0.06%	6.09	1.121	8.84%	6.09	0.000
14	-0.01%	5.90	1.028	8.82%	5.90	0.000

Notes: The table reports the average abnormal returns (*AAR*) and the cumulative average abnormal returns (*CAAR*) over the event window. In the baseline analysis, the event window spans from $\tau_1=-2$ to $\tau_2=14$ and lasts seventeen seconds ($L_2=17$), whereas the estimation window spans from $\tau_0=-10$ to $\tau_1=-3$ and lasts eight seconds ($L_1=8$).

Table A.3
Event study results (robustness check).

Panel A: set events ($N=365$)						
Time (τ)	AAR	t-statistic (θ_1)	p-value	CAAR	t-statistic (θ_2)	p-value
-2	0.09%	0.29	0.775	0.09%	0.29	0.775
-1	0.19%	0.57	0.567	0.28%	0.61	0.544
0	0.12%	0.38	0.700	0.40%	0.72	0.473
1	0.23%	0.72	0.474	0.63%	0.98	0.327
2	0.41%	1.26	0.209	1.04%	1.44	0.150
3	0.96%	2.96	0.003	2.00%	2.52	0.012
4	0.30%	0.92	0.359	2.30%	2.68	0.007
5	0.40%	1.24	0.215	2.70%	2.94	0.003
6	1.40%	4.32	0.000	4.10%	4.22	0.000
7	0.30%	0.94	0.349	4.40%	4.30	0.000
8	-0.04%	-0.11	1.088	4.36%	4.06	0.000
9	-0.05%	-0.15	1.120	4.32%	3.85	0.000
10	0.15%	0.47	0.641	4.47%	3.82	0.000
11	-0.04%	-0.12	1.094	4.43%	3.65	0.000
12	0.12%	0.37	0.713	4.55%	3.62	0.000
13	-0.11%	-0.33	1.260	4.44%	3.43	0.001
14	-0.17%	-0.52	1.395	4.27%	3.20	0.001
15	1.06%	3.26	0.001	5.33%	3.88	0.000
16	-0.09%	-0.29	1.225	5.24%	3.71	0.000
17	-0.11%	-0.33	1.255	5.13%	3.54	0.000

Panel B: tie-break events ($N=79$)						
Time (τ)	AAR	t-statistic (θ_1)	p-value	CAAR	t-statistic (θ_2)	p-value
-2	0.06%	0.16	0.876	0.06%	0.16	0.876
-1	0.83%	2.04	0.042	0.89%	1.55	0.121
0	0.62%	1.54	0.124	1.51%	2.15	0.031
1	1.13%	2.78	0.006	2.64%	3.25	0.001
2	0.99%	2.45	0.014	3.63%	4.01	0.000
3	3.38%	8.32	0.000	7.01%	7.05	0.000
4	0.86%	2.12	0.034	7.87%	7.33	0.000
5	0.89%	2.20	0.028	8.76%	7.64	0.000
6	0.83%	2.05	0.041	9.59%	7.88	0.000
7	1.43%	3.53	0.000	11.03%	8.59	0.000
8	-0.05%	-0.13	1.103	10.97%	8.16	0.000
9	-0.07%	-0.17	1.135	10.91%	7.76	0.000
10	1.08%	2.67	0.008	11.99%	8.19	0.000
11	-0.07%	-0.17	1.135	11.92%	7.85	0.000
12	-0.06%	-0.15	1.122	11.86%	7.55	0.000
13	0.14%	0.35	0.726	12.00%	7.39	0.000
14	0.18%	0.45	0.649	12.18%	7.28	0.000
15	0.15%	12.19	0.000	12.33%	7.95	0.000
16	-0.14%	-0.34	1.267	12.19%	7.61	0.000
17	-0.10%	-0.24	1.188	12.09%	7.31	0.000

Notes: The table reports the average abnormal returns (AAR) and the cumulative average abnormal returns (CAAR) over the event window. In the robustness analysis, the event window spans from $\tau_1=-2$ to $\tau_2=17$ and lasts 19 seconds ($L_2 = 19$), whereas the estimation window spans from $\tau_0=-17$ to $\tau_1=-3$ and lasts 15 seconds ($L_1=15$).

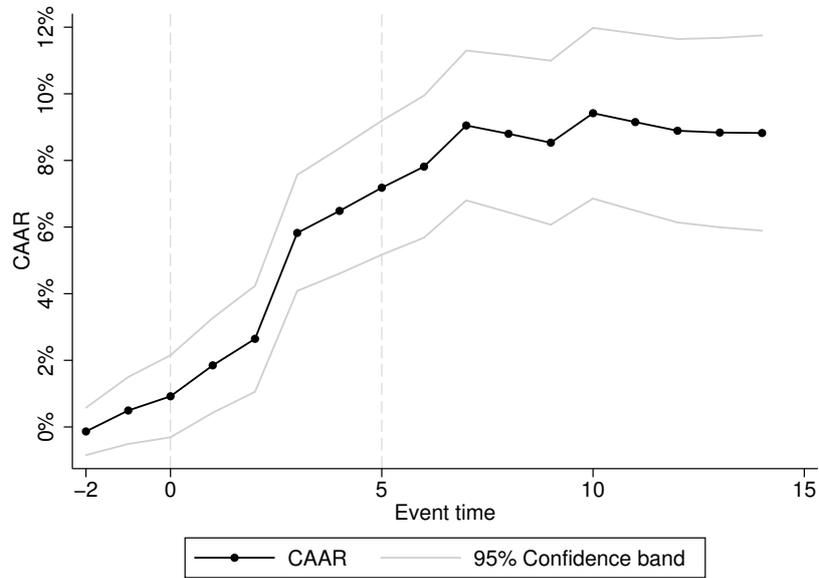


Figure A.1
Cumulative average abnormal returns (*CAAR*) for the tie-break events.

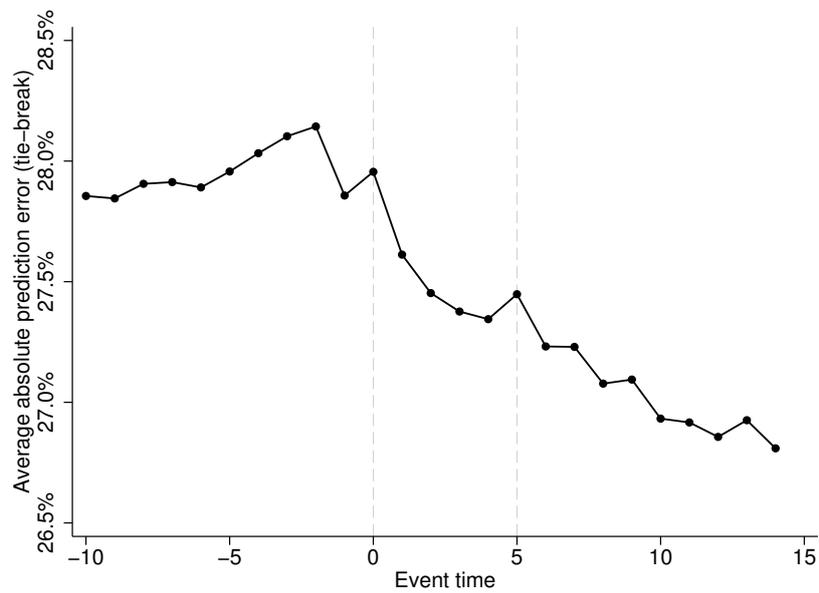


Figure A.2
Mean absolute prediction error (*APE*) for the tie-break events.

References

- Betfair. (2015). *Annual report*. Retrieved 2017-06-10, from www.paddypowerbetfair.com/~media/Files/P/Paddy-Power-Betfair/archive/annual-report-general-meeting/betfair/annual-report-2015.pdf
- Betfair. (2017). *Terms and conditions*. Retrieved 2017-07-07, from <http://www.betfair.com/aboutUs/Terms.and.Conditions/#terms-betting-general>
- Brown, A. (2012). Evidence of in-play insider trading on a UK betting exchange. *Applied Economics*, *44*(9), 1169–1175.
- Brown, A., & Yang, F. (2017). Slowing down fast traders: Evidence from the Betfair speed bump (forthcoming). *Journal of Financial Markets*.
- Camerer, C. F. (1989). Does the basketball market believe in the hot hand? *The American Economic Review*, *79*(5), 1257–1261.
- Carlton, D. W., & Fischel, D. R. (1983). The regulation of insider trading. *Stanford Law Review*, 857–895.
- Chakravarty, S., & McConnell, J. J. (1997). An analysis of prices, bid/ask spreads, and bid and ask depths surrounding Ivan Boesky's illegal trading in Carnation's stock. *Financial Management*, 18–34.
- Chakravarty, S., & McConnell, J. J. (1999). Does insider trading really move stock prices? *Journal of Financial and Quantitative Analysis*, *34*(2), 191–209.
- 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*(Part B), 614–629.
- Cornell, B., & Sirri, E. R. (1992). The reaction of investors and stock prices to insider trading. *The Journal of Finance*, *47*(3), 1031–1059.
- Cox, S. (2015). *Why tennis courtsiding was my dream job*. Retrieved 2017-06-16, from www.bbc.com/news/magazine-32402945
- Croxson, K., & Reade, J. (2014). Information and efficiency: Goal arrival in soccer betting. *The Economic Journal*, *124*(575), 62–91.
- Dickson, C. (2015). Courtsiding in sport: Cheating, sharp practice or merely irritating? *Sports Law eJournal*.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, *25*(2), 383–417.
- Fama, E. F. (1991). Efficient capital markets 2. *The Journal of Finance*, *46*(5), 1575–1617.
- Fernandes, N., & Ferreira, M. A. (2009). Insider trading laws and stock price informa-

- tiveness. *Review of Financial Studies*, 22(5), 1845–1887.
- Fishman, M. J., & Hagerty, K. M. (1992). Insider trading and the efficiency of stock prices. *The RAND Journal of Economics*, 23(1), 106–122.
- Franck, E., Verbeek, E., & Nüesch, S. (2010). Prediction accuracy of different market structures: bookmakers versus a betting exchange. *International Journal of Forecasting*, 26(3), 448–459.
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance*, 46(1), 179–207.
- Hutchins, B. (2014). *Game, Set, Cash! Inside the secret world of international tennis trading*. Nero.
- Ikenberry, D., Lakonishok, J., & Vermaelen, T. (1995). Market underreaction to open market share repurchases. *Journal of Financial Economics*, 39(2), 181–208.
- Jaffe, J. F. (1974). Special information and insider trading. *The Journal of Business*, 47(3), 410–428.
- Ke, B., Huddart, S., & Petroni, K. (2003). What insiders know about future earnings and how they use it: Evidence from insider trades. *Journal of Accounting and Economics*, 35(3), 315–346.
- Keown, A. J., & Pinkerton, J. M. (1981). Merger announcements and insider trading activity: An empirical investigation. *The Journal of Finance*, 36(4), 855–869.
- Kooij, W. J., Stokking, H. M., van Brandenburg, R., & de Boer, P.-T. (2014). Payout delay of TV signals: Measurement system design, validation and results. In *Proceedings of the 2014 ACM international conference on interactive experiences for TV and online video* (pp. 23–30).
- Lakonishok, J., & Lee, I. (2001). Are insider trades informative? *Review of Financial Studies*, 14(1), 79–111.
- Landi, G. (2015). *Betting exchange: La rivoluzione del trading sportivo*.
- Leland, H. E. (1992). Insider trading: Should it be prohibited? *Journal of Political Economy*, 100(4), 859–887.
- MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1), 13–39.
- Manne, H. G. (1966). Defense of insider trading. *Harvard Business Review*, 44(6), 113–122.
- Meulbroeck, L. K. (1992). An empirical analysis of illegal insider trading. *The Journal of Finance*, 47(5), 1661–1699.
- Mitchell, M., Pulvino, T., & Stafford, E. (2004). Price pressure around mergers. *The Journal of Finance*, 59(1), 31–63.

- Rozeff, M. S., & Zaman, M. A. (1988). Market efficiency and insider trading: New evidence. *Journal of Business*, *61*(1), 25–44.
- Schnytzer, A., & Shilony, Y. (1995). Inside information in a betting market. *The Economic Journal*, *105*(431), 963–971.
- Spann, M., & Skiera, B. (2003). Internet-based virtual stock markets for business forecasting. *Management Science*, *49*(10), 1310–1326.
- Townend, M. (2016). Tennis gambling market second only to football for bookmakers after boom in online and in-play betting. *Daily Mail Online*.
- Vaughan Williams, L. (2009). *Information efficiency in financial and betting markets*. Cambridge University Press.
- Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, *18*(2), 107–126.