Adverse Selection, Speed Bumps and Asset Market Quality*

Alasdair Brown†
University of East Anglia

Fuyu Yang‡
University of East Anglia

April 28, 2015

*We would like to thank Ben Mcquillin, Subhasish Modak Chowdhury, Arnold Polanski, James Watson, and an audience at UEA for comments. Alasdair Brown acknowledges financial support from the British Academy and Leverhulme Trust, as part of the project 'The Role of Trading Volume in Asset Mispricing' (Ref: SG140097). The research presented in this paper was carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia. The usual disclaimer applies.

†Corresponding Author. School of Economics, University of East Anglia, Norwich NR4 7TJ, U.K.. Email: alasdair.brown@uea.ac.uk

‡School of Economics, University of East Anglia, Norwich NR4 7TJ, U.K.. Email: fuyu.yang@uea.ac.uk

Abstract

Recent evidence suggests that the fastest algorithmic traders in financial markets profit at the expense of slower traders. One solution gaining traction is a 'speed-bump', which introduces a delay between the time in which an order is submitted, and when it is processed. We conduct an impact evaluation of the speed bump's effectiveness on Betfair, a betting exchange, where this design has been in force for more than a decade. We find that increases in the duration of the delay led to improvements in liquidity (measured by bid-ask spreads and depth) and market quality (measured by order frequency and volume).

JEL Classification: G10, G12, G14

Keywords: algorithmic trading, high-frequency trading, speed-bump, market liquidity

1 Introduction

Financial market traders compete over differences in the speed of their trading. These differences are measured in milliseconds, microseconds, and even nanoseconds. The fastest algorithmic, high-frequency traders (HFTs) obtain, process, and trade upon financial news before others - often using automated trading strategies themselves - have had a chance to react and update their quotes. Baron *et al.* (2014) find that the fastest HFTs trading the E-mini S&P 500 futures contract are consistently the most profitable. Moreover, these traders make approximately 45% of their profits from adversely-selecting slightly slower liquidity-providing HFTs. As shown by the model of Biais *et al.* (2014), the rents available to fast traders lead to over-investment in speed-enhancing technologies. The most famous example of this over-investment is the high-speed fibre-optic cable - built by Spread Networks in 2010 at a cost of $300m - which shaved just 1.4 milliseconds off the time of round-trip communication between Chicago and New York. The competitive advantage from using this new cable was brief, as microwave towers brought the time of round-trip communication down by a further 4.1 milliseconds in 2012.[1]

[1]A new cable crossing the Atlantic in order to bring down communication times between London and New York, dubbed the Hibernia Express, will be in service from summer 2015. For more details on these

Early work on the effect of algorithmic trading on market liquidity (e.g. Hendershott *et al.* (2011)) had actually found gains from greater automation. Automation can reduce the costs of attention, and allow market-makers to manage their inventory more effectively, which should reduce the size of bid-ask spreads and other measures of trading costs. However, while some automation may be beneficial, there is growing evidence that increases in the efficiency of automation (increases in trader speed) have led to lower liquidity. Menkveld and Zoican (2014) find that increased speed of trade execution on the Nasdaq OMX in Scandinavia (of the order of 2.25 milliseconds) had a detrimental effect on liquidity (the effective spread increased by 32%). In addition, Chaboud *et al.* (2014) find that the introduction of algorithmic trading in the foreign exchange market imposed higher adverse selection costs on slower traders. In other words, it would appear that the fastest traders use their advantage to pick off their slower counterparties, rather than provide liquidity more efficiently themselves.

What can be done to end this arms-race in speed technologies? What can be done to protect liquidity-providing traders from the fastest traders? Is there a market design that will harness the benefits of algorithmic trading, whilst reducing or eliminating the incentives to be the fastest?

One possible solution is a 'speed bump' or 'order processing delay' of the type utilised by IEX, of 'Flash Boys', Lewis (2014) fame, and soon to be adopted by Toronto Stock Exchange (TSX) Alpha. The speed bump works by adding a delay between the time in which an order is submitted, and the time in which the order is logged on the exchange. The delay proposed by TSX Alpha is between 5 and 25 milliseconds. The idea is that if a trader has a fleeting informational advantage - related to a very recent announcement, for example - this advantage will have dissipated by the time that their order has been executed. In protecting other traders from this algorithmic form of adverse selection, the speed bump may reduce spreads and other measures of trading costs. Moreover, by taking away the profitability of extremely low-latency trading strategies, such a market design could eliminate the incentive to be the fastest, and put an end to the arms-race in speed-enhancing technologies.

While the use of a speed bump is a recent innovation in financial markets, a similar

trading-motivated developments, see 'Raging Bulls: How Wall Street Got Addicted to Light-Speed Trading, Jerry Adler, *Wired*, 8th March 2012.

market design has been used in U.K. betting markets for more than a decade. In this paper, we conduct an impact evaluation of the effectiveness of the speed bump in improving market liquidity and quality on Betfair, the world's largest betting exchange.

The Betfair betting exchange is a standard limit order book for the trading of bets on soccer, tennis and other major sports. As on financial exchanges of the same type, traders providing liquidity submit limit orders (which sit in the book until an offsetting order arrives), and traders consuming liquidity submit market orders (which execute at prices already quoted in the book). The exchange facilitates trading both prior to sporting events, and during events ('inplay'). A common problem in inplay betting is that a subset of market order traders, present in the stadia, may adversely select the stale quotes of limit order traders. Due to inevitable delays in the broadcast of sporting events, such traders - dubbed 'court-siders' in tennis and 'pitch-siders' in soccer - may know the outcome of the last point or be aware of a recent goal. These traders are analogous to the fastest HFTs, who trade quickest on the most recent macroeconomic and firm-specific news (as in Brogaard *et al.* (2014)). Depending on the severity of the adverse selection problem, liquidity providers may widen the spread - to recoup any losses to pitch/court-siders with gains from other bettors - or even cease providing liquidity altogether (Glosten and Milgrom (1985)).

Betfair use a speed bump to counter this problem. Any order submitted inplay is delayed by 5-9 seconds before it is logged on the exchange. (There is no delay prior to play). This means that those providing liquidity through existing limit orders have time in which to cancel their quotes if information arrives in that 5-9 second window. (Cancellations are not subject to the speed bump). Therefore, even if a subset of traders receive information faster, by the time they are permitted to execute their trade, the information is also in the hands of all other traders.

To evaluate the success, or otherwise, of the speed bump, we need an appropriate counterfactual: how liquid would the betting exchange have been without the order processing delay? Unfortunately, the delay has been in place since the beginning of inplay trading, which means that there is no simple comparison to be made. We therefore exploit variation in the length of the delay over five seasons of English Premier League soccer. In 2008/9, the majority of matches were subject to a 5 second delay. (The delay within each match is constant). In each subsequent season, the proportion of matches subject to a longer delay of 8 or 9 seconds was steadily increased until 91% of matches had order processing delays of 8 or 9 seconds in 2012/13. With a longer delay, the scope for pitch-siders to bet on goals that have already occurred reduces. (Pitch-siders would need to be aware of a goal at least 8 or 9 seconds quicker than their counterparty, as opposed to only 5 seconds quicker with the shorter delay). Therefore, if the speed bump is an effective tool, we would expect that market liquidity would improve each season as the proportion of matches with longer delays increases. Importantly, we can observe liquidity prior to matches, when there is no adverse selection from pitch-siders, and no speed bump delay. This allows us to control for any time-trends in overall liquidity on the exchange over the five seasons.

We find that increases in the average order processing delay led to an improvement in a range of market liquidity and quality measures. Over the five seasons, inplay quoted spreads dropped by 7.3%, effective spreads fell by 3.8%, and inplay quoted depth at the best three prices (on both sides of the book) increased by 19%. Perhaps more striking, we find that the number and size of inplay orders (our primary measures of market quality) increased season-by-season. The frequency of inplay orders increased by 87%, and inplay order volumes increased by 35%. In short, we find that all of our measures of market liquidity and quality improved as the order processing delay was increased. It would appear that increases in the delay protected the majority of bettors from being picked off by pitch-side bettors, and in turn trading costs decreased, and volumes increased.

While our main analysis accounts for time-trends in liquidity on the exchange using pre-match liquidity as a benchmark (in a difference-in-difference setting), it is possible that we could be confounding the effect of the speed bump with general increases in interest in inplay betting. To allay such concerns, we runs a series of placebo tests using Wimbledon tennis betting data from the same exchange. Unlike with soccer matches, Betfair kept the inplay tennis delay to 5 seconds throughout the same period. If we are confounding the effect of the soccer speed bump with a steady increase in inplay betting popularity, we would then expect to see similar improvements in tennis inplay market liquidity over the same period. Yet, we find no significant improvements in spreads or depth for tennis betting. Only inplay trading volumes increased over the same period, but there was no contemporaneous increase in the frequency of orders. We therefore conclude that the improvements in market liquidity and quality that we observe in soccer inplay betting are, in the main, due to increases in the order processing delay.

This paper contributes to a nascent literature on market designs to mitigate the speed advantages of the fastest HFTs. Budish *et al.* (2013, 2014) propose the use of frequent batch auctions conducted at one second intervals. Traders would submit limit orders as before - specifying the price and volume at which they wish to trade. Every second, these orders would be aggregated by the exchange, demand and supply schedules would be drawn, and market clearing prices and quantities would be determined. Any orders not fulfilled in a batch auction - if the trader bid less, or asked more, than the clearing price - would be sent to the next batch auction, one second later, to try again. Importantly, within each batch auction, no priority would be given to the time at which the order was submitted.

An alternative solution, suggested by Gai *et al.* (2013), is to de-regulate the minimum tick size in financial markets. These authors show that HFT participation is greater in stocks where the minimum tick size is binding (i.e. where spreads cannot get smaller). Furthermore, they find that HFT activity is associated with higher levels of quote cancellations, more short-term price volatility, and little depth (quoted volume) at the best price. They propose further reductions in the minimum tick size to allow liquidity-providing HFTs to compete, once again, on price rather than speed.

In this paper we investigate an alternative solution to the problem of speed advantages in asset markets. Uniquely amongst the solutions offered, the speed bump has been at work on an exchange for more than a decade. This allows us to carry out a detailed evaluation study - using quasi-experimental methods - of its effect on market liquidity and quality. Furthermore, the speed bump appears to be an appealing solution to financial exchange operators, given than IEX and TSX Alpha implemented such a design feature without a known impact evaluation of its effectiveness.

This brings us to the external validity of our results. Many of the same features found in financial markets can be found in this betting market. Trade takes place on a open limit order book, the most common form of financial exchange (Parlour and Seppi (2008)), and bettors can execute trades manually or algorithmically through the Betfair Application

Program Interface (API). While financial market volumes may dwarf those found in betting markets, the matches we study are subject to more than 9 billion GBP worth of bets. This scale ensures that participants are incentivised to price assets efficiently, and provide liquidity competitively.

Testifying to this external validity, there is a rich history of research using betting data to consider questions related to financial markets. In particular, betting markets have proved useful for questions of market efficiency, as the fundamental value of bets are revealed at the end of proceedings. Many authors have focused on pari-mutuel markets, including Ali (1977), Snyder (1978) and Snowberg and Wolfers (2010). In a pari-mutuel market, there is no intermediary, and the payout from bets is inversely proportional to the amount wagered. Others authors - including Shin (1993) and Vaughan Williams and Paton (1997) - have looked at U.K. bookmaker markets, where there is an intermediary setting prices. More recently, focus has turned to the betting exchanges, otherwise labelled 'prediction markets'. Tetlock (2008), Croxson and Reade (2014), Brown (2014), Brown and Yang (2014) and Choi and Hui (2014) all look at the efficiency of these markets. Fewer papers have examined the liquidity of these markets. Brown (2012) detected the presence of court-siders - using data on bid-ask spreads - in the Wimbledon Final of 2008. Our aim in this current paper is to look at a measure intended to counter such adverse selection. In doing so, we take advantage of a market design innovation (the speed bump) that hit betting markets first, but has subsequently been adopted by a number of financial markets.

The rest of this paper is structured as follows. In Section 2 we describe the betting exchange, the adverse selection problem posed by pitch-siders, and the implementation of the speed bump. In Section 3 we describe the data and our measures of market liquidity and quality. Section 4 contains the analysis, and Section 5 concludes.

2 Inplay Betting and the Speed Bump

Betfair is the largest betting exchange in the world. Founded in the U.K. in 2000, the exchange facilitates the trading of bets on a range of sporting, political, and celebrity events. The exchange operates as a standard limit order book, with the option of long (back) or short (lay) positions, and the option of limit or market orders. The development of this exchange - and its now defunct competitor Flutter - represented a departure from the bookmaker dominated system in the U.K.. On the exchanges, bettors could provide liquidity themselves, and more easily take short positions.[2]

An illustration of the website view of the Betfair limit order book can be found in the top panel of Figure 1. The screenshot is taken from the match between Crystal Palace and Manchester City on the 6th April 2015 (a match outside of our sample). Odds are displayed with associated depth underneath. Back odds on the left-hand-side allow a punter to bet on a team, and lay odds on the right-hand-side allow for betting against a team. A back quote of 7, for example, offers a punter 6 GBP (plus the return of their stake) for each 1 GBP that they place on that outcome. A lay quote of 7.2 - on the other side of the book - requests a payout of 6.2 plus the stake. Punters can choose to trade at current quotes (with a market order), or provide quotes themselves (with a limit order). The screenshot shows the limit orders in the book at the time.

While Betfair originally facilitated only pre-event sports betting, the development of inplay betting was also pioneered on the exchange. Here, punters could wager on an event as it unfolded. The inception of this type of betting was accompanied by a new form of adverse selection. Due to inevitable delays in television broadcasts, punters present at a sporting event had a fleeting informational advantage. Those present could be aware of a goal in soccer, the outcome of a point in tennis, or a fall in a horse race, a few seconds before those watching the action at home.

The presence of such court-siders, as they were dubbed in tennis, or pitch-siders, as they were dubbed in soccer, received a great deal of media attention in January 2014, when a court-sider was arrested at the Australian Open tennis championships in Melbourne.[3] This individual was inputting scores into his phone, which would then be fed into an algorithmic

[2]Prior to the development of exchanges, a punter looking to bet with a bookmaker against a given horse, for example, would have to bet on all other competitors in the race.

[3]See, for example, 'Tennis's New Concern: Data Harvesting', Greg Bishop and John Martin, *New York Times*, January 21st 2014, 'Tennis Arrest at Company Set Up by Former Betfair Staff', Alistair Osborne, *The Telegraph*, 21st January 2014, and 'Inside the Shadowy World of High-Speed Tennis Betting', Carl Bialik, *538*, May 29th 2014.

trading strategy implemented on Betfair. By consistently being a point ahead of other bettors watching the action at home in the U.K., courtsiders could expect to profit from hundreds of fleeting informational advantages.[4]

The advantage available to pitch-siders or court-siders depends on the lag in television broadcasts. The length of broadcast delays depend on the technology used both by the broadcaster and the viewer. Kooij (2014) conducted a study of broadcast delays, with the aim of facilitating the synchronisation of different broadcasts for interactive programmes. He estimates that the minimum delay from filming to receiving a broadcast is 4 seconds. This delay increases by approximately 2 seconds if the broadcast is via satellite (rather than terrestrial), increases further for high-definition transmissions, and an additional delay is incorporated if the transmission must be passed between foreign and domestic broadcasters. Even longer lags are found in internet streams.

In order to counter the type of adverse selection that arises due to broadcast delays, Betfair have implemented a speed bump, or order processing delay, since the inception of inplay trading. The duration of the order processing delay does, however, vary significantly across sports. Horse racing is subject to a delay of just 1 second, which is clearly insufficient to protect bettors at home. This perhaps explains why Brown and Yang (2014) found that speculative trade, by market order traders, was highly predictive of fundamentals during races. For tennis the delay is 5 seconds. The presence of court-siders at the Australian Open suggests that this is insufficient, at least for that particular broadcast. Finally, soccer is subject to a delay of 5-9 seconds, with higher delays in more recent seasons. The calculations of Kooij (2014) suggest that that delays at the upper end (8 or 9 seconds) may, just, be enough to eradicate a pitch-sider's advantage.

The speed bump is illustrated in the bottom panel of Figure 1. The screenshot is taken from the same match as earlier, but this time a bet has been placed. The ticker in the top-left hand corner illustrates how much of the order processing delay is still left to run (in this case 1 second). Once that 1 second is up, the order will be logged in the book, and execute with

[4] Although the individual was arrested under suspicion of harming tennis's integrity, the practice of court-siding has little to do with match-fixing or corruption. Court-siding likely has effects only on the equity, and liquidity, of associated betting markets.

a limit order if such a limit order has not been cancelled.

In this study we use the steady increases in the order processing delay for soccer matches, season by season, as our 'treatment'. We then compare the drift in liquidity and market quality over the seasons for inplay bets - subject to pitch-siding and increases in the order processing delay - with any drift that might be observed in pre-match trading, where there is neither this type of adverse selection, nor a speed bump in operation. To ensure that we are not confounding the effect of increases in the order processing delay with changes in the interest in inplay trading, we also run a series of placebo tests using Wimbledon tennis data. Over the same number of years, Betfair kept the tennis order processing delay fixed at 5 seconds.

Why did Betfair raise the order processing delay for soccer betting, but not for tennis betting? Unfortunately, we can only speculate about the reasons. Certain articles written around the time of the arrest of the court-sider in Melbourne emphasised that the firm behind the court-sider was run by former Betfair employees.[5] A second, perhaps related, point is that just as financial exchanges are conflicted in their attitudes to the most predatory high-frequency traders - they generate order flow (and commissions), but are likely to reduce liquidity - betting exchanges undoubtedly face the same dilemma. A third possible explanation is that pitch-siding in soccer creates more obvious losers. If a bettor has their quotes picked-off by a pitch-sider who knew that a goal has already been scored, the unfortunate bettor will undoubtedly know, afterwards, that they had been at an informational disadvantage. On the other hand, a bettor who is continually one point behind the court-sider may be oblivious to this fact, as one point has less impact on the price than a goal. This factor may, or may not, explain Betfair's willingness to lengthen the order processing delay for soccer bets, but not for tennis bets.

Before we proceed, it is worth explaining why we do not use variation, within season, in the assignment of short and long order processing delays. Why, for example, do we not compare a match in 2009/10 which was subject to a 5 second delay, with a match in the same season with a 8 second delay? The reason is that the assignment of the length of order

[5] See, for example, 'Tennis Arrest at Company Set Up by Former Betfair Staff', Alistair Osborne, *The Telegraph*, 21st January 2014.

processing delay is non-random within seasons. Matches played at 3pm on a Saturday - and therefore not legally televised - are disproportionately more likely to be assigned an 8 or 9 second order processing delay compared to televised matches played at other times. We will illustrate this point in more detail in the next section. It is the case, however, that - for both Saturday 3pm matches and televised matches played at other times - the probability of being 'treated' with a longer order processing delay increases each season. In other words, a match with any given set of characteristics has an increased chance of getting a longer order processing delay in a later season than in an earlier season.

3 Data

We collected data on English Premier League soccer betting, on Betfair, for the seasons 2008/9 through to 2012/13. This data was acquired from Fracsoft, a third-party provider of Betfair data. The Fracsoft data includes second-by-second snapshots of the limit order book, from the morning of each match, through to its conclusion. Incomplete data was available for the 2006/7 and 2007/8 seasons, but this largely related to the marquee televised games - which undoubtedly receive more attention and volume - and therefore due to selection concerns we left these seasons out of our analysis. The five seasons worth of data that we use relate to 5,211 separate bets, which, as each match has 3 bets (win/lose/draw), means that we have data, on average, on 347 out of a possible 380 matches a season. This yields more than 108 million observations of the Betfair limit order book.

For our placebo tests we collected six years of Betfair betting data on the Men's Wimbledon Tennis Championships from 2008 to 2013 inclusive. This data comes from the same source (Fracsoft) as the soccer data and comes in the same form. As with our main data, the limit order book is sampled from the morning of each match until its conclusion. There are potentially 768 matches (128 per year) for the six years of the Championships we study. However, there are withdrawals, byes and missing data, and therefore we have data on 641 separate matches. This yields approximately 35.4 million observations of the Betfair limit order book.

Our first comments concern the duration of the order processing delay. As shown in

Table 1, the average inplay order processing delay, for our main soccer data, increased each season from 2008/9 to 2012/13. In the first season the average delay was 6.3 seconds, and in the last season the delay was 7.65 seconds. The smallest increase in the average delay was between 2008/9 and 2009/10, when the delay increased by 21 milliseconds. As described in the Introduction, the delay within each match is fixed, which means that the number of matches with longer delays of 8 or 9 seconds increases season-by-season. In contrast, the inplay delay for Wimbledon tennis betting stayed constant at 5 seconds throughout the sample period. For both soccer and tennis, the order processing delay was 0 seconds prior to the start of play.

We also summarise the information related to the main soccer data in Figure 2. In the left panel we display the pre-match data, with the inplay data in the right panel. Anticipating some of our later analysis, we also break the sample down into Saturday 3pm matches, and non-Saturday 3pm matches. The former cannot be legally televised (but can often be watched using illegal internet streams), and the majority of the latter are legally broadcast on satellite television. As illustrated in this figure, Saturday 3pm matches have higher average delays than the televised matches. For both types of matches, however - both televised and non-televised - the average inplay order processing delay broadly increases season-by-season.

We are interested, firstly, in the effect of changes in the order processing delay on market liquidity. For this purpose, we use the following four measures of liquidity:

- Quotes: an indicator variable equalling 1 if there are, in that second, quoted prices on both sides of the book (bid and ask/back and lay).

An extreme indication of adverse selection would be the absence of any liquidity provision at all. In the Glosten and Milgrom (1985) model, liquidity would disappear in such a manner if all market order traders were informed (about the outcome of the match, and the fundamental value of the bets).

- Quoted spread: defined, in implied win probability terms, as $(1/B - 1/L) * 100$, where B are the best back odds, and L are the best lay odds.

Odds quoted on Betfair include the stake, and therefore odds of 3, will return 2 GBP plus the 1 GBP stake in the case of a win. Therefore, suppose back odds are 4, and lay odds are 5. This would result in a (very large) spread of $(1/4 - 1/5) * 100 = 5$. More generally, a widening of the quoted spread indicates greater adverse selection (see Glosten and Milgrom (1985) again).

- Effective spread: defined, in implied win probability terms, as $100 * 2 * abs(1/M_t - (1/B_{t-1} + 1/L_{t-1})/2)$, where B are the best back odds, L are the best lay odds, M is the last transaction price (in odds), and abs(.) is the absolute difference.

The effective spread takes the absolute difference between the last transaction price, and the midpoint of the bid-ask spread in the previous second. The term is multiplied by 2 to make it equivalent to the quoted spread, otherwise it would be a half spread. The effective spread captures the costs that traders actually pay to trade. On the one hand, the effective spread may be smaller than the quoted spread, as traders are more likely to trade when quoted spreads are smaller. On the other hand, there are instances when effective spreads can be larger than quoted spreads. Trades of a large size may not be solely executed at the best quotes, and will need to walk down the book to the 2nd best quotes, the 3rd best quotes, and so on. Effective spread therefore captures this fact, and incorporates both the quoted spread, and to a certain extent, the quoted depth.

- Depth: defined as the sum of the volumes associated with the best three quotes, on both the back and lay side of the book.

Liquidity providers can also adjust the volume they are willing to trade in the presence of informed traders (see Dupont (2000)). In the presence of adverse selection, we would expect to see lower depth quoted at the best prices.

The aim of the speed bump is that it reduces the adverse selection problem for liquidity providers. If the informed's informational advantage is fleeting, then it should have disappeared by the time they are able to cross the speed bump and execute their trades. However, while the speed bump may lessen adverse selection, it could have negative consequences for the overall quality of asset markets. For example, such a market design may make it more difficult to execute trades, particularly of a significant size (as liquidity providers will have cancelled their quotes before the order is logged on the exchange).

With this in mind, we are also interested in the following two measures of asset market quality:

- Order: an indicator variable equalling 1 if there was an order in the last second.

The Fracsoft data includes information on the total matched volume each second, for each bet. Therefore any increase in this indicates that a market order has been executed. If the speed bump is detrimental to trade, this measure should decline, inplay, over the seasons.

- Volume: the size of the order, in GBP, conditional on there being an order.

This information can also be calculated from the total matched volume measure provided by Fracsoft. If the speed bump is detrimental to trade, we might expect this measure to decline, inplay, season-by-season. Summary statistics on the four liquidity measures and the two market quality measures can be found in Table 2.

4 Analysis

For the initial part of our analysis, we estimate an equation of the following form:

$$Outcome_{it} = \beta_0 + \beta_1 Season_{it} + \beta_2 Inplay_{it} + \beta_3 Season_{it} Inplay_{it} + e_{it} \qquad (1)$$

$Outcome_{it}$ is one of our market liquidity or quality measures, $Season_{it}$ is the number of the season running from 2008/9 (number 1) to 2012/13 (number 5), $Inplay_{it}$ is an indicator equalling 1 if the match in question has begun, $Season_{it}Inplay_{it}$ is an interaction between the two, and e_{it} is an error term. i is a subscript to indicate the bet in question (e.g. Manchester United to beat Arsenal on the 3rd of October at Old Trafford), and t is a subscript to indicate the particular second that the bet prices are sampled (i.e. 3rd October 16.21.03 pm).

The idea behind this difference-in-difference specification is as follows. The $Season_{it}$ will capture general drifts in the $Outcome_{it}$ over time. The $Inplay_{it}$ indicator will capture the declines in liquidity, and increases in orders, that occur as a match begins. Our key variable of interest is the interaction term, $Season_{it}Inplay_{it}$, which will capture the specific change in the $Outcome_{it}$, inplay, over time, as the average order processing delay increases.

We begin, in Table 3, with an analysis of the Quotes indicator, which equals 1 if there are quotes to both back and lay the bet in that second. (For this analysis, and all the following regressions, we cluster standard errors at the bet-level). To accompany the results in Table 3, we also present averages of this outcome variable in Figure 3. Pre-match averages for each season are in the left panel, and inplay averages are in the right panel. At this stage, we are only interested in the full sample, displayed in green. Quotes are almost always available prior to matches, and are available more than 90% of the time inplay. Our regression results indicate that there is a slight increase in the probability of there being quotes inplay over the seasons. In other words, increases in the order processing delay are accompanied by improvements in this coarse liquidity measure. It should be noted, however, in the context of our difference-in-difference estimates, that the probability of pre-match quotes can only go down (from 1), while inplay this measure has some freedom to tick up or down over the seasons.

We turn next, in Table 3 and Figures 4 and 5, to quoted spreads and effective spreads. We find significant decreases in both quoted and effective spreads, inplay, over the five seasons. This is set against a backdrop of slight, albeit insignificant, increases in pre-match quoted and effective spreads over the same time-frame. In terms of the raw data, average inplay quoted spreads dropped by 7.3% from 2008/9 to 2012/13, and average inplay effective spreads dropped by 3.8%. In other words, the cost of trading seems to have decreased as the order processing delay increased.

Our final liquidity measure is depth, which is the sum of quoted volume at the three best back and lay quotes. These results are in Table 3, with graphics in Figure 6. Against a backdrop of declining depth in the pre-match markets, we find increases in inplay depth. From 2008/9 to 2012/13, average inplay depth increased by 19%, while pre-match depth decreased by 38%. To sum up at this stage, we have evidence to suggest that all four liquidity measures improved as a result of increases in the order processing delay.

One concern with speed bumps is that they may make it more difficult to execute trades. Market order traders are subject to the delay, but those who have posted limit orders are free to cancel them without delay. This may mean that orders are more often cancelled than executed-with, and that trade frequency and volume declines. There is little point having ostensibly liquid markets, if no-one actually gets to trade.

With this in mind, in Table 3 and Figures 7 and 8, we examine the effect of increases in the order processing delay on the frequency of executed-orders and the size of these orders. As illustrated quite clearly in Figure 6, the frequency of inplay orders increases at a faster rate, over the seasons, than the frequency of pre-match orders. Inplay order frequency almost doubled from 2008/9 to 2012/13. Similarly, the average volume of each of these orders increased. Against a backdrop of declining order size in pre-match markets, the size of inplay orders increased by 35% from 2008/9 to 2012/13. It would appear that rather than prohibiting the execution of trade, the increases in the speed bump delay encouraged more, and larger, bets.

Perhaps the strongest concern with our analysis thus far, is that we may be confounding the effect of the speed bump with changes in the popularity of inplay betting. The argument may go that pre-match betting interest was static or declining, but inplay betting had a surge in popularity which led to greater liquidity, more orders and larger volumes. Our approach to this concern is to replicate our analysis with an analysis of Wimbledon tennis betting on Betfair. If inplay betting popularity surged relative to pre-match betting, we should observe a similar effect in this sport. (Of course, we cannot rule out the possibility that there was a football-specific surge in inplay betting popularity, but we have no reason to believe that there was). Importantly, Betfair kept the inplay order processing delay for tennis fixed at 5 seconds throughout the same period.

In Table 4 we display the results of our placebo tests on Wimbledon betting data. Season is replaced by Year - which runs from 2008 (number 1) to 2013 (number 6) - but otherwise the regressions are identical. In terms of our four liquidity measures, we find a slight decline in the probability of quotes (our coarsest measure of liquidity) and no significant effect on our three other measures. In other words, there is nothing to suggest that there were similar improvements in inplay liquidity in the tennis betting market. In terms of market quality, we do find that the volume of inplay tennis orders increased significantly. However, this is not accompanied by an increase in the frequency of orders, as it was for the soccer data. In short, our placebo tests lead us to believe that the improvements in inplay soccer liquidity, and to a lesser extent market quality, were not due to general increases in inplay betting interest, but instead were due to increases in the order processing delay.

In our final piece of analysis, we return to the main soccer data. There is reason to believe that there may be differences between matches that are televised, and those that are not. In the U.K., it is illegal to broadcast Saturday 3pm matches. The rationale is that this will harm attendances for lower league matches. Therefore, to watch these Saturday 3pm matches, punters would either need to use an illegal satellite dish (to watch the broadcast of another country), or follow a foreign broadcast on the internet. As discussed in Section 2, internet streams tend to have the longest delays. The specific lag, of course, depends on the foreign broadcaster and the internet technology. It is possible, however, that the protection afforded to bettors from the speed bump is more important for these Saturday 3pm matches, as they are either not watching the match, or are watching it many seconds behind the pitch-sider.

It is with this idea in mind, that we proceed to our analysis in Table 5. We regressed our six outcomes variables on the three variables from Table 3, but this time added an indicator for Saturday 3pm matches, and every possible combination of interactions between the Saturday 3pm indicator, the Season, and the Inplay indicator. Markets on Saturday 3pm matches tend to be less liquid and attract less orders, particularly inplay (due to the difficulties in television access). Our key variables of interest, however, are the Season-Inplay indicator from before, and the Saturday 3pm-Season-Inplay indicator added to these latest specifications. The former informs us of whether there was a baseline effect on televised matches from increases in the order processing delay, and the latter tells us whether any such effect was larger or smaller for Saturday 3pm matches.

Our results are as follows. We find that the probability of finding quotes was improved by the greater speed bump (with no differences across televised/non-televised matches), and quoted and effective spreads improved, particularly for Saturday 3pm matches. If we consider spreads to be our preferred measure of the cost of trading, our results suggest that the effect of increases in the speed bump delay was concentrated in matches where punters either did not watch the match, or were at the greatest disadvantage relative to pitch-siders. This, in our view, supports the case for the effectiveness of the speed bump. Having said that, lower trading costs do not necessarily go hand-in-hand with greater market quality. Depths, order frequency and order volume improved for both types of matches, but more so for televised matches.

5 Conclusion

Financial markets are now dominated by computerised trading, where orders are determined by algorithms. This has created a technological arms-race, where the profitability of trading firms is determined by the speed of their trading (Baron *et al.* (2014)). By being the fastest, firms can trade on information (e.g. macroeconomic/firm news, or the contents of the limit order book) before others have had a chance to react. Such inequity in financial markets runs the risk of reducing participation rates and market liquidity.

Profiting from advance access to information is, of course, nothing new. Since the first financial markets, traders have attempted to gain, and exploit, informational advantages over other traders. Nathan Rothschild reportedly heard of Napoleon's 1815 defeat at Waterloo through his family's network of couriers, and used this information to purchase stock in victorious London.[6] The difference with current practices is two-fold. Firstly, informational advantages are acquired, and utilised, many times each day, and therefore have the potential to continually impair market liquidity. The Rothschilds did not trade at such high-frequency. A second difference is that the solution to this new form of adverse selection may lie in market design rather than regulation (e.g. insider trading laws). It is the search for an effective market design to mitigate the advantages to the fastest traders that motivates this paper.

We evaluate the effectiveness of a speed bump, where there is a delay between the time in which an order is submitted, and the time that it is logged on an exchange. If a fast algorithmic trader in financial markets, or a pitch-sider in betting markets, has a fleeting informational advantage, such an advantage should have dissipated before their order can pick off a stale limit order provided by another trader. In protecting liquidity providers from this form of adverse selection, the speed bump has the potential to reduce the cost of trading,

[6] While the purchase of stock may have been profitable, it is likely that Napoleon's defeat was, overall, bad news for the Rothschilds. Ferguson (1999) states that the Rothschilds had accumulated vast amounts of bullion to be lent to the Duke of Wellington to pay troops during what was expected to be a protracted war. The relatively swift defeat of Napoleon left the Rothschilds holding more than they wished.

and increase market participation rates.

We find that increases in the length of the order processing delay, on Betfair, a U.K. betting exchange, led to improvements in liquidity, as measured by bid-ask spreads and quoted depths available at the best quotes. Furthermore, we find that greater protection afforded to traders from the threat of pitch-siders, led to more frequent trade, and larger volume involved in each trade. In short, market liquidity and quality responded well to the more stringent use of the speed bump. While we cannot say whether the speed bump reduces expenditure in speed-enhancing technologies (co-location and faster cables in financial markets, attendance at matches in betting markets) such an enquiry is a natural next step for this research.

References

- Ali, M., M., (1977). Probability and Utility Estimates for Racetrack Bettors. *Journal of Political Economy*, 85, 803-815.

- Baron, M., Brogaard, J., Kirilenko, A., (2014). Risk and Return in High Frequency Trading. Working paper, http://ssrn.com/abstract=2433118.

- Biais, B., Foucault, T., Moinas, S., (2014). Equilibrium Fast Trading. *Journal of Financial Economics*, forthcoming.

- Brogaard, J., Hendershott, T., Riordan, R., (2014). High Frequency Trading and Price Discovery. *Review of Financial Studies*, 27, 2267-2306.

- Brown, A., (2012). Evidence of In-Play Insider Trading on a U.K. Betting Exchange. *Applied Economics*, 44, 1169-1175.

- Brown, A., (2014). Information Processing Constraints and Asset Mispricing. *Economic Journal*, 124, 245-268.

- Brown, A., Yang, F., (2014). The Role of Speculative Trade in Market Efficiency: Evidence from a Betting Exchange. Applied and Financial Economics working paper series, No. 68.

- Budish, E., Cramton, P., Shim, J., (2013). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. Working paper.
- Budish, E., Cramton, P., Shim, J., (2014). Implementation Details for Frequent Batch Auctions: Slowing Down Markets to the Blink of an Eye. *American Economic Review*, P&P, 104, 418-424.
- Chaboud, A., P., Chiquoine, B., Hjalmarsson, E., Vega, C., (2014). Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *Journal of Finance*, 69, 2045-2084.
- Choi, D., Hui, S., K., (2014). The Role of Surprise: Understanding Over- and Underreactions Using In-play Soccer Betting. *Journal of Economic Behavior and Organization*, 107, 614-629.
- Croxson, K., Reade, J., J., (2014). Information and Efficiency: Goal Arrival in Soccer Betting. *Economic Journal*, 124, 62-91.
- Dupont, D., (2000). Market-Making, Prices, and Quantity Limits. *Review of Financial Studies*, 13, 1129-1151.
- Ferguson, N., (1999). *The House of Rothschild: Money's Prophets 1798-1848*. Penguin.
- Gai, J., Yao, C., Ye, M., (2013). The Externalities of High-Frequency Trading. Working paper, http://ssrn.com/abstract=2066839.
- Glosten, L., R., Milgrom, P., R., (1985). Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics*, 14, 71-100.
- Hendershott, T., Jones, C., M., Menkveld, A., J., (2011). Does Algorithmic Trading Improve Liquidity? *Journal of Finance*, 66, 1-33.
- Kooij, W., (2014). Playout Delay of TV Broadcasting. University of Twente.
- Lewis, M., (2014). *Flash Boys: Cracking the Money Code*. Allen Lane.

- Menkveld, A., J., Zoican, M., A., (2014). Need for Speed? Exchange Latency and Market Quality. Working paper.

- Parlour, C., A., Seppi, D., J., (2008). Limit Order Markets: A Survey. *Handbook of Financial Intermediation and Banking*, edited by Boot, A., W., A., Thakor, A., V., Elsevier.

- Shin, H., S., (1993). Measuring the Incidence of Insider Trading in a Market for State-Contingent Claims. *The Economic Journal*, 103, 1141-1153.

- Snowberg, E., Wolfers, J., (2010). Explaining the Favorite-Long Shot Bias: Is it Risk Love or Misperceptions? *Journal of Political Economy*, 118, 723-746.

- Snyder, W., W., (1978). Horse Racing: Testing the Efficient Markets Model. *Journal of Finance*, 33, 1109-1118.

- Tetlock, P., C., (2008). Liquidity and Prediction Market Efficiency. *mimeo*.

- Vaughan Williams, L., Paton, D., (1997). Why is there a Favourite-Longshot Bias in British Racetrack Betting Markets? *The Economic Journal*, 107, 150-158.

Figures and Tables

(a) Limit Order Book

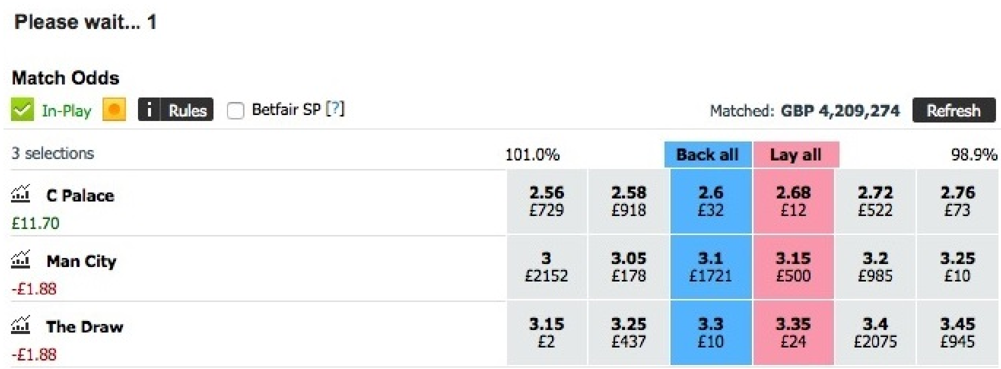

(b) Speed Bump in Operation

Figure 1: **Betfair Limit Order Book and Speed Bump**. Screenshots of the Betfair limit order book, as accessed through the website rather than the Application Program Interface (API). These screenshots are taken, inplay, from the match between Crystal Palace and Manchester City on the 6th April 2015 (outside our data-set). In panel (a) no bet has been placed, while in panel (b) a bet has been placed and the order processing delay is ticking down (see top left).

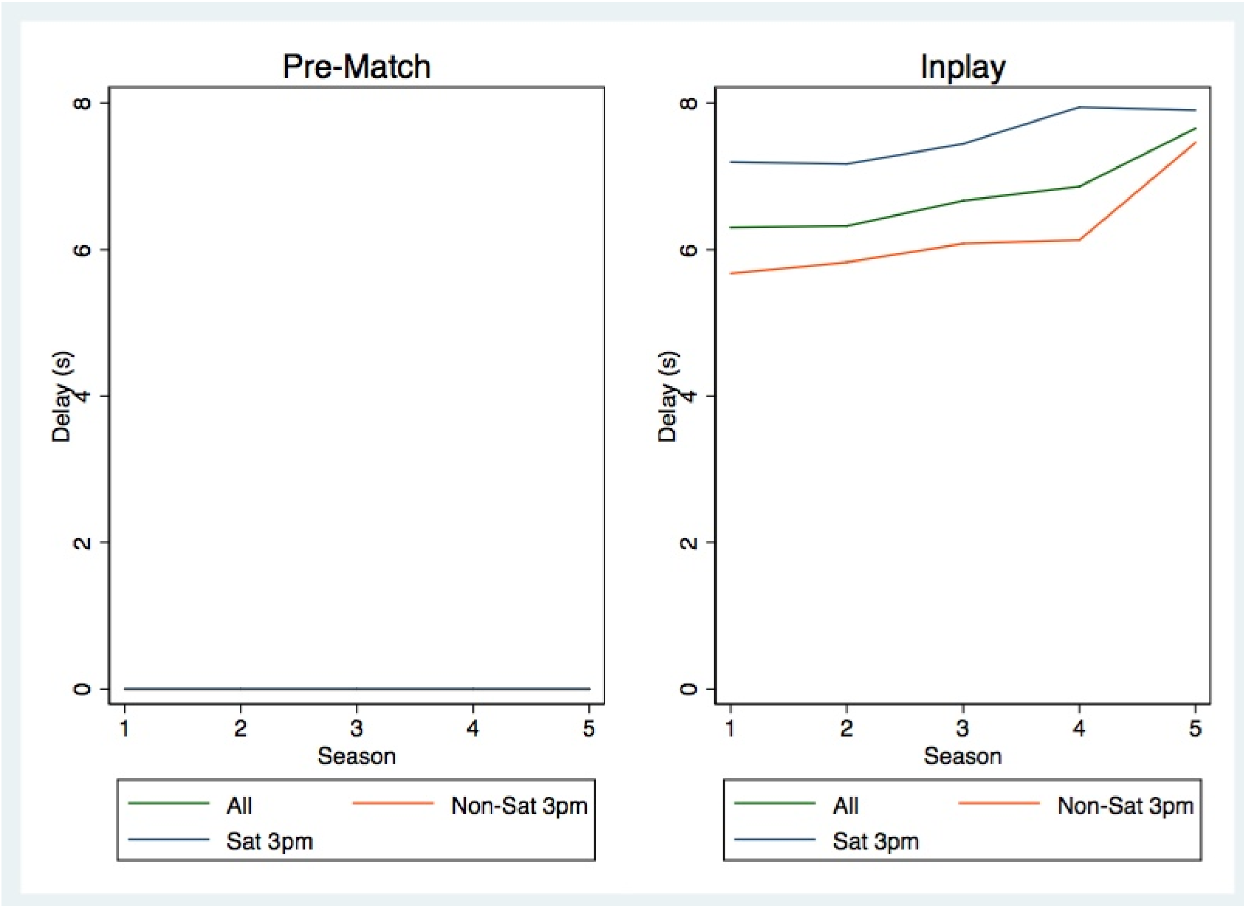

Figure 2: **Speed Bump Delays**. A comparison of average order processing delays, in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used. Delays are measured in seconds.

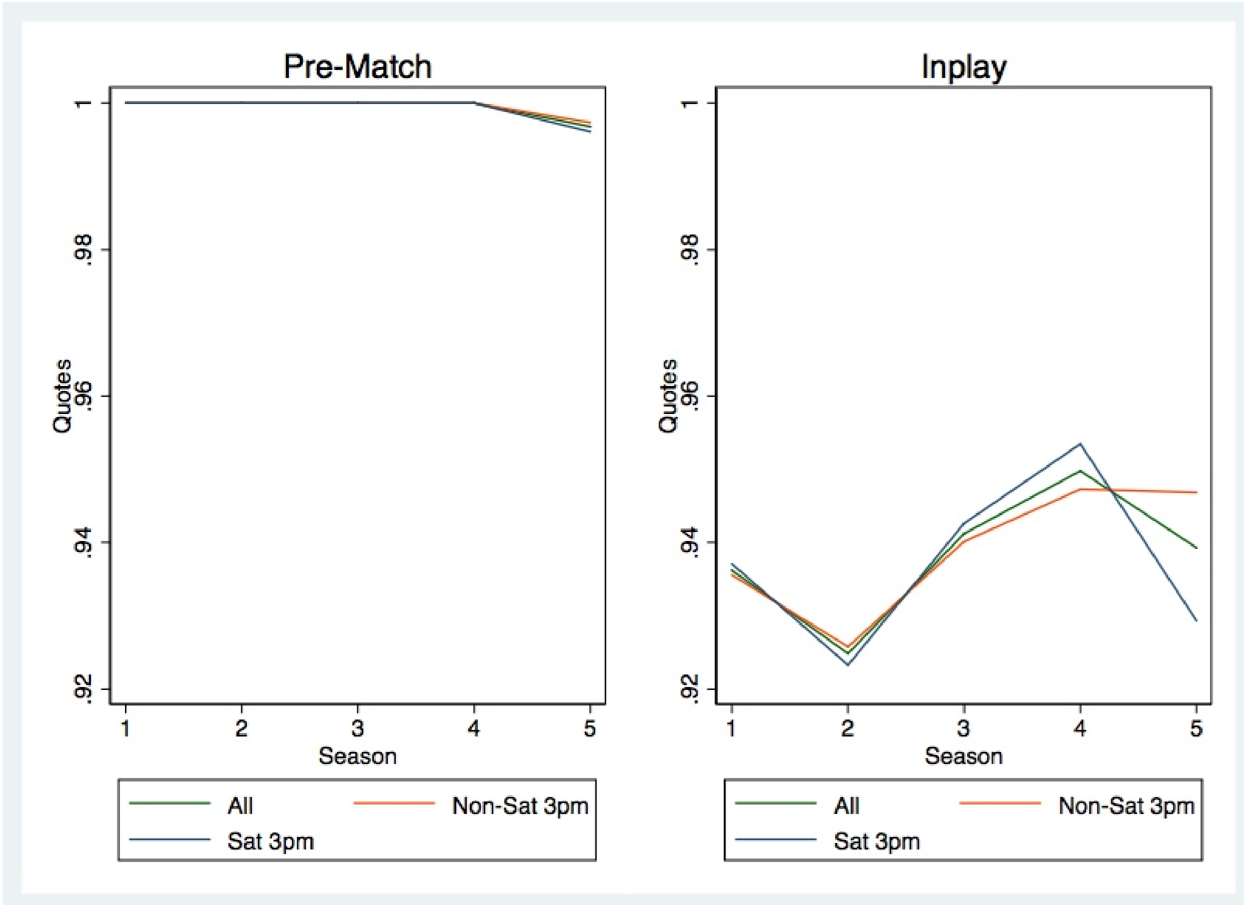

Figure 3: **Quotes**. A comparison of the proportion of seconds where there are quotes to both back and lay a bet, in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

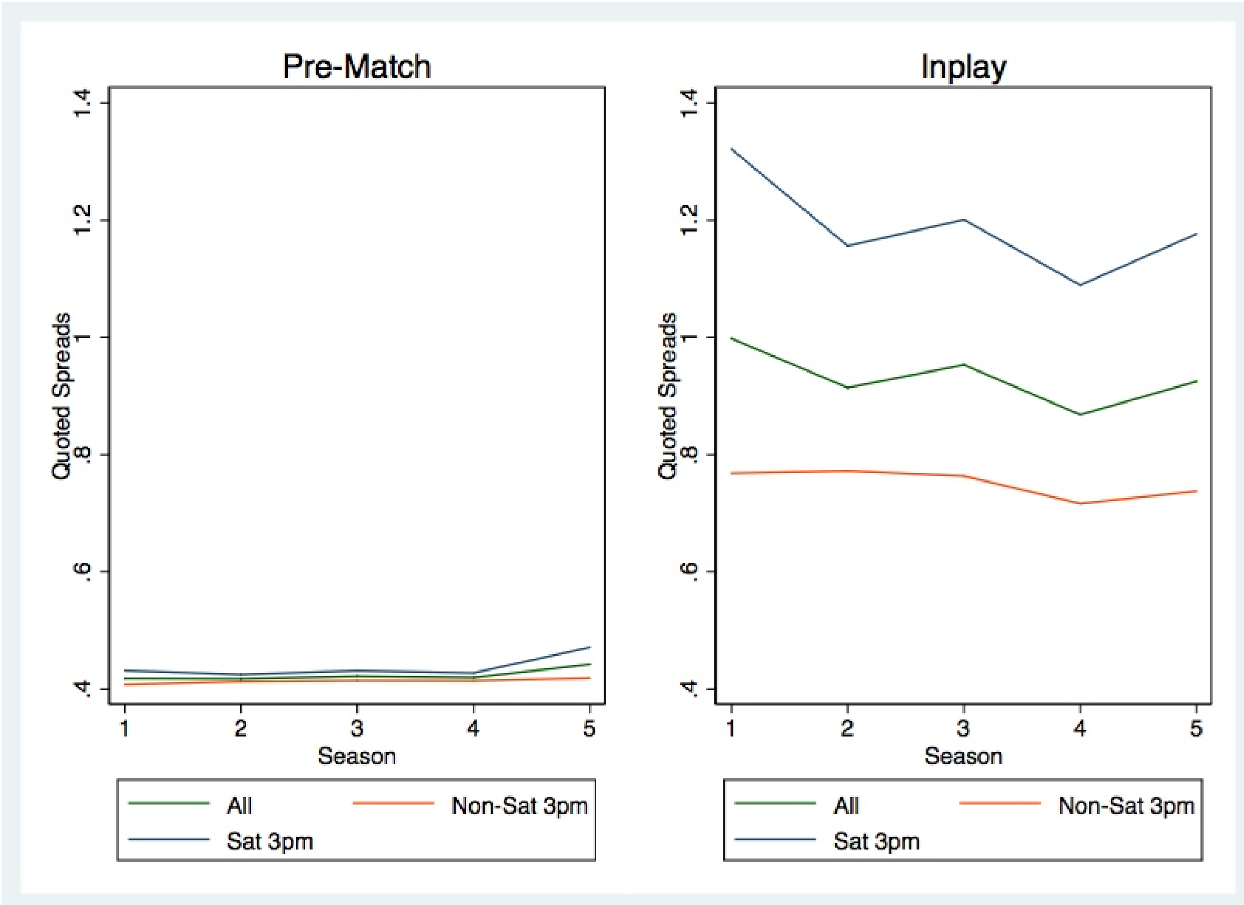

Figure 4: **Quoted Spreads**. A comparison of average quoted back-lay spreads (as defined in Section 3), in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

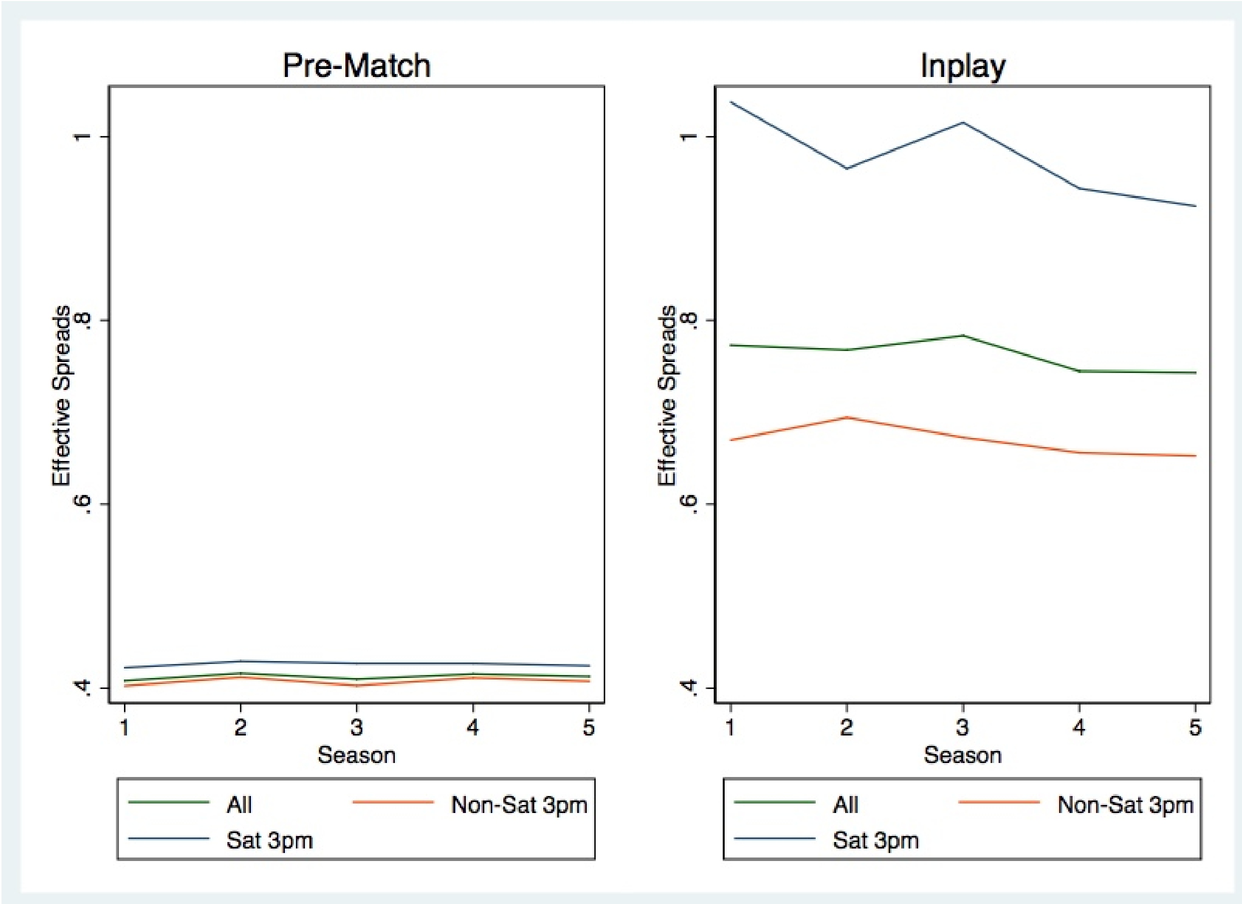

Figure 5: **Effective Spreads**. A comparison of average effective back-lay spreads (as defined in Section 3), in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

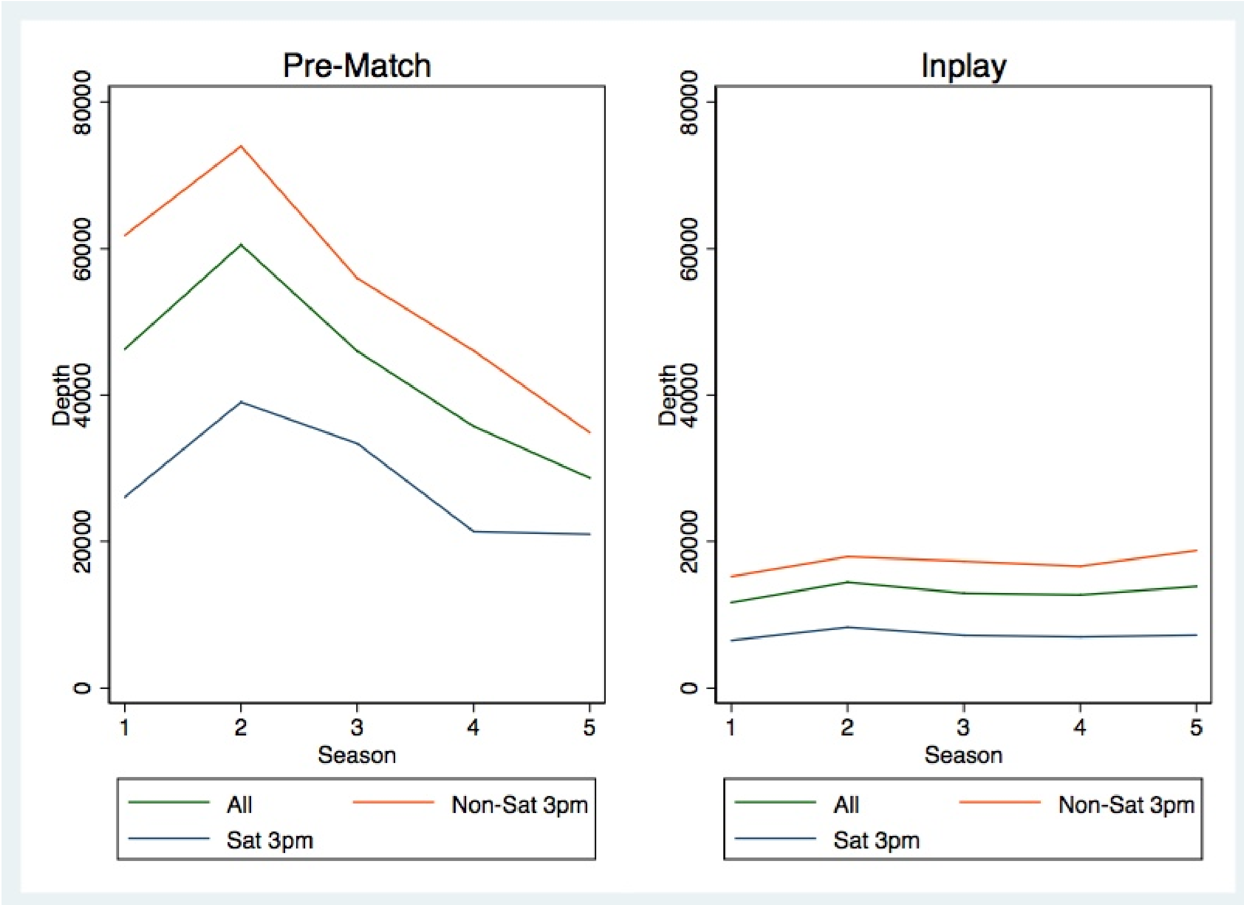

Figure 6: **Depths**. A comparison of average quoted depth at the best 3 back and lay prices, in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

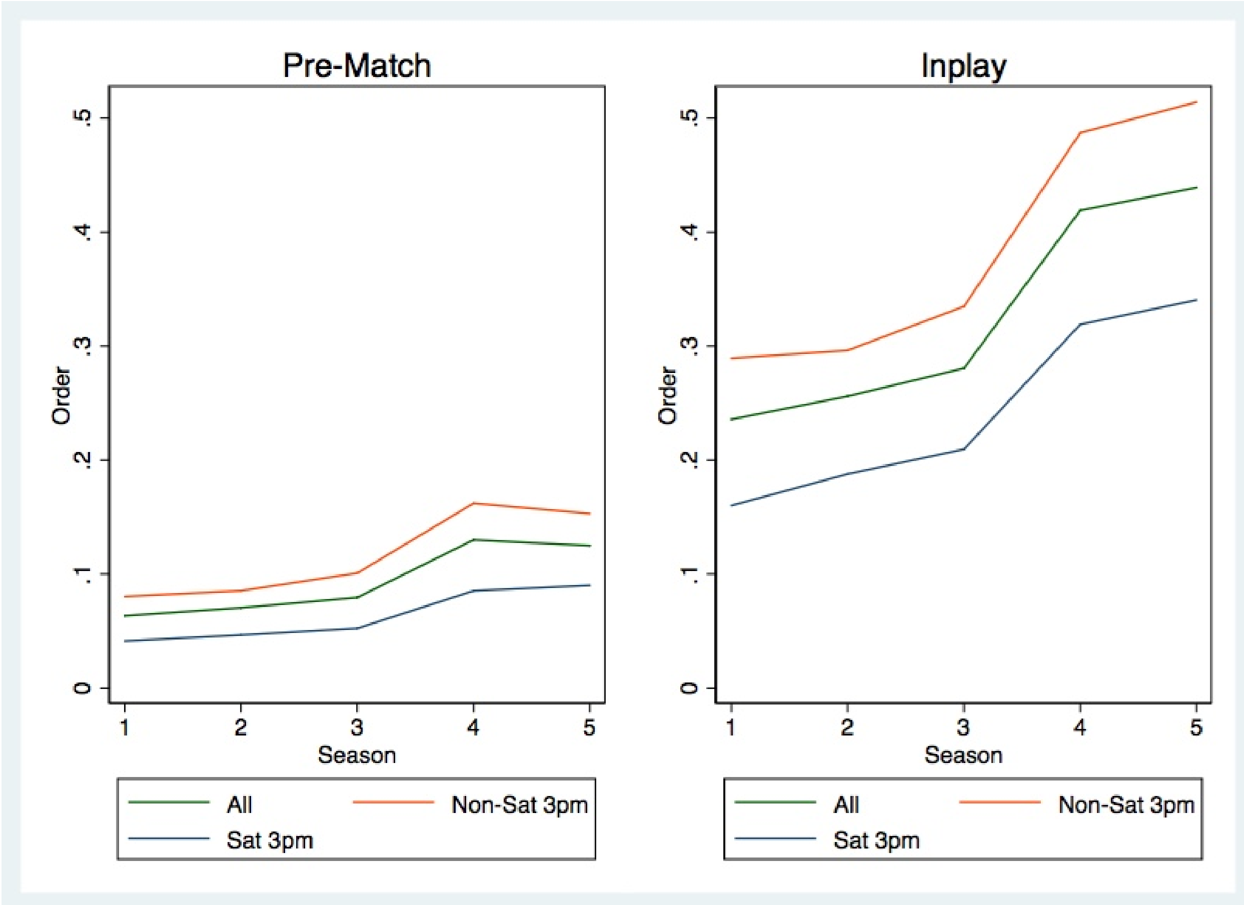

Figure 7: **Orders**. A comparison of the proportion of seconds where there was an order, in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

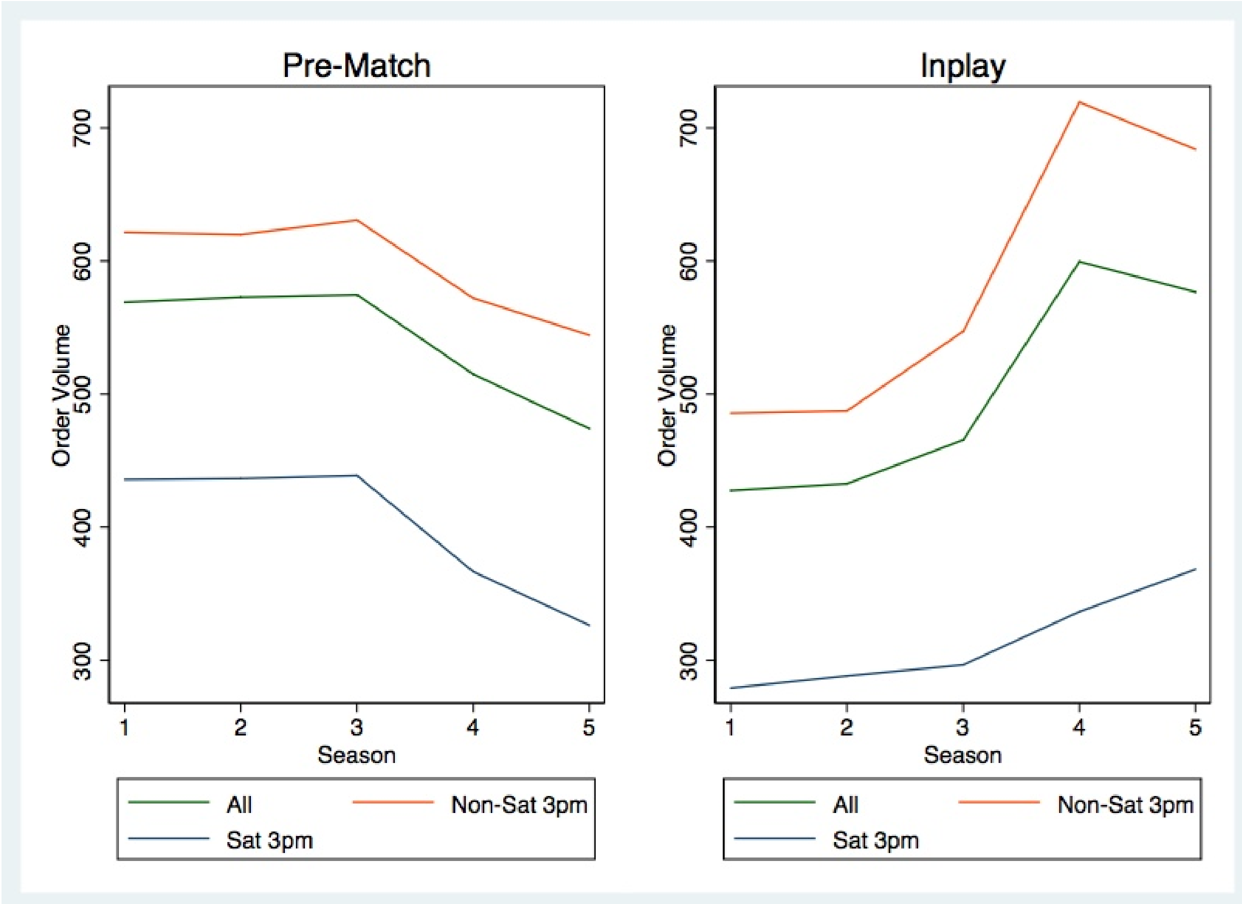

Figure 8: **Volumes**. A comparison of average trade size (in GBP), in the pre-match period (left panel) and the inplay period (right panel). The full data-set is displayed in green, with Saturday 3pm and non-Saturday 3pm sub-samples displayed in blue and orange respectively. Data from 5 seasons of Betfair betting on the English Premier League from 2008/9 (Season 1) to 2012/13 (Season 5) are used.

Table 1: Speed Bump Delays					
Main Data: English Premier League					
Pre-Match	Obs	Mean	Std. Dev.	Min	Max
All Seasons	73,463,253	0	0	0	0
In-Play	Obs	Mean	Std. Dev.	Min	Max
2008-9	7,261,470	6.300	1.147	5	8
2009-10	7,822,380	6.321	1.182	5	8
2010-11	7,574,457	6.669	1.323	5	9
2011-12	7,281,597	6.862	1.457	5	9
2012-13	7,662,105	7.655	1.002	5	9
Placebo Data: Wimbledon					
Pre-Match	Obs	Mean	Std. Dev.	Min	Max
All Years	23,356,038	0	0	0	0
Inplay	Obs	Mean	Std. Dev.	Min	Max
All Years	12,307,584	5	0	5	5

Speed Bump Delays. Summary statistics on the duration of the order processing delay (measured in seconds). Data is broken down into pre-match and inplay periods, and English Premier League (top panel) and Wimbledon (bottom panel) betting data. The soccer data covers the 2008/9 season to the 2012/13 season, while the tennis data covers the tournament every year from 2008-2013.

Table 2: Summary Statistics					
Main Data: English Premier League					
Outcome Variable	Obs	Mean	Std. Dev.	Min	Max
Quotes	108,477,531	.97	.141	0	1
Quoted Spread	106,262,585	.580	1.902	.001	98.9
Effective Spread	18,142,462	.627	1.92	0	186.2
Depth	105,010,329	34,199	99,021	12	5,064,690
Order	108,477,531	.169	.374	0	1
Volume	18,142,462	518	2,902	.007	786,540
Placebo Data: Wimbledon					
Outcome Variable	Obs	Mean	Std. Dev.	Min	Max
Quotes	35,467,938	.96	.195	0	1
Quoted Spread	34,059,702	1.8	4.39	.001	90.4
Effective Spread	2,372,264	1.56	2.95	0	179.1
Depth	30,953,873	53,313	201,810	12	5,113,224
Order	35,467,938	.068	.251	0	1
Volume	2,372,264	933	4,848	.007	744,745

Summary Statistics. Summary statistics on each of the outcome variables described in Section 3. Quotes is an indicator variable equalling 1 if there are quotes to both back and lay a bet in that second, Quoted Spread is defined in Section 3, Effective Spread is defined in Section 3, Depth is the available volume quoted at the best 3 back and lay prices, Order is an indicator variable equalling 1 if an order was executed in that second, and Volume is the trade size per second conditional on an order. Data is broken down into English Premier League (top panel) and Wimbledon (bottom panel) betting data. The soccer data covers the 2008/9 season to the 2012/13 season, while the tennis data covers the tournament every year from 2008-2013.

Table 3: Difference-in-Difference (Main Data)			
Outcome:	Quotes	Quoted Spread	Effective Spread
Intercept	1.00***	0.408***	0.410***
	(.000)	(.007)	(.003)
Season	-0.000.	0.005	0.000
	(.000)	(.003)	(.001)
Inplay	-0.072***	0.58***	0.378***
	(.003)	(.021)	(.01)
Season*Inplay	0.003***	-0.024**	-0.009**
	(.001)	(.006)	(.003)
No. of Clusters (Bets)	5,211	5,211	5,211
No. of Observations	108,477,531	106,262,585	18,142,462
R^2	0.041	0.015	0.007
Outcome:	Depth	Order	Volume
Intercept	62,003***	0.038***	620***
	(3,901)	(.001)	(21.4)
Season	-6,151***	0.018***	-26.3***
	(954)	(.000)	(6.02)
Inplay	-49,639***	0.115***	-260***
	(3,474)	(.003)	(18.4)
Season*Inplay	6,408***	0.038***	73.3***
	(858)	(.001)	(5.75)
No. of Clusters (Bets)	5,208	5,211	5,211
No. of Observations	105,010,329	108,477,531	18,142,462
R^2	0.024	0.101	0.0004

Difference-in-Difference (Main Data: English Premier League). Regressions of 6 outcome variables on the season (1-5) that the match took place, an inplay indicator, and an interaction between the season and the inplay indicator. The interaction term captures the 'treatment' effect of increases in the order processing delay from season to season. The 6 outcomes variables are as follows. Quotes is an indicator variable equalling 1 if there are quotes to both back and lay a bet in that second, Quoted Spread is defined in Section 3, Effective Spread is defined in Section 3, Depth is the available volume quoted at the best 3 back and lay prices, Order is an indicator variable equalling 1 if an order was executed in that second, and Volume is the trade size per second conditional on an order. Standard errors (in parentheses) are clustered at the bet-level, and ***, **, * and . indicates significance at the 0.1%, 1%, 5% and 10% level respectively.

Table 4: Difference-in-Difference (Placebo Data)			
Outcome:	Quotes	Quoted Spread	Effective Spread
Intercept	0.999***	0.679***	0.565***
	(.000)	(.02)	(.021)
Year	0.000	-0.018**	-0.008
	(.000)	(.005)	(.005)
Inplay	-0.096***	3.55***	1.16***
	(.009)	(.263)	(.092)
Year*Inplay	-0.005*	0.06	0.017
	(.002)	(.067)	(.023)
No. of Clusters (Bets)	1,261	1,261	1,261
No. of Observations	35,467,938	34,059,702	2,372,264
R^2	0.08	0.158	0.024
Outcome:	Depth	Order	Volume
Intercept	47,059***	0.014***	490.2***
	(13,891)	(.001)	(55.8)
Year	4,625	0.000	24.6
	(3,815)	(.000)	(16.2)
Inplay	-36,116**	0.136***	152.4.
	(11,369)	(.01)	(86.2)
Year*Inplay	368	0.003	77.9*
	(2,939)	(.002)	(27.5)
No. of Clusters (Bets)	1,261	1,261	1,261
No. of Observations	30,953,873	35,467,938	2,372,264
R^2	0.007	0.08	0.002

Difference-in-Difference (Placebo Data: Wimbledon Tennis Championships). Regressions of 6 outcome variables on the year (1-6) that the match took place, an inplay indicator, and an interaction between the year and the inplay indicator. The interaction term captures the 'treatment' effect of increases in the order processing delay from year to year. The 6 outcomes variables are as follows. Quotes is an indicator variable equalling 1 if there are quotes to both back and lay a bet in that second, Quoted Spread is defined in Section 3, Effective Spread is defined in Section 3, Depth is the available volume quoted at the best 3 back and lay prices, Order is an indicator variable equalling 1 if an order was executed in that second, and Volume is the trade size per second conditional on an order. Standard errors (in parentheses) are clustered at the bet-level, and ***, **, * and . indicates significance at the 0.1%, 1%, 5% and 10% level respectively.

Table 5: Triple Difference (Main Data)			
Outcome:	Quotes	Quoted Spread	Effective Spread
Intercept	1.00***	0.406***	0.404***
	(.001)	(.004)	(.004)
Season	-0.000	0.002.	0.000
	(.000)	(.001)	(.001)
Inplay	-0.075***	0.381***	0.288***
	(.004)	(.02)	(.009)
Season*Inplay	0.005***	-0.014*	-0.008**
	(.001)	(.005)	(.002)
Saturday 3pm	0.000	0.005	0.02**
	(.001)	(.015)	(.008)
Saturday 3pm*Season	-0.000	0.006	-0.000
	(.000)	(.007)	(.002)
Saturday 3pm*Inplay	0.007	0.503***	0.340***
	(.006)	(.045)	(.024)
Saturday 3pm*Season*Inplay	-0.003	-0.029*	-0.016*
	(.002)	(.014)	(.007)
No. of Clusters (Bets)	5,211	5,211	5,211
No. of Observations	108,477,531	106,262,585	18,142,462
R^2	0.0413	0.0193	0.0108
Outcome:	Depth	Order	Volume
Intercept	80,106***	0.048***	661.9***
	(5,694)	(.002)	(27.1)
Season	-8,423***	0.022***	-21.4**
	(1,385)	(.001)	(7.77)
Inplay	-64,575***	0.140***	-266.9***
	(5,085)	(.005)	(23.9)
Season*Inplay	8,980***	0.042***	84.8***
	(1,258)	(.001)	(7.71)
Saturday 3pm	-43,592***	-0.026***	-168.6***
	(7,449)	(.003)	(38.6)
Saturday 3pm*Season	5,628**	-0.008**	-9.49
	(1,830)	(.001)	(10.6)
Saturday 3pm*Inplay	35,247***	-0.067***	16.4
	(6,660)	(.006)	(33.7)
Saturday 3pm*Season*Inplay	-6,175***	-0.006**	-30.3**
	(1,649)	(.002)	(10.4)
No. of Clusters (Bets)	5,208	5,211	5,211
No. of Observations	105,010,329	108,477,531	18,142,462
R^2	0.0389	0.1168	0.0021

Difference-in-Difference-in-Difference (Main Data: English Premier League). Regressions of 6 outcome variables on the season (1-5) that the match took place, an inplay indicator, an interaction between the season and the inplay indicator, an indicator for Saturday 3pm matches, an interaction between the Saturday 3pm indicator and the season, an interaction between the Saturday 3pm indicator and the inplay indicator, and an interaction between the Saturday 3pm indicator, the season and the inplay indicator. The interaction term between the season and the inplay indicator captures the 'treatment' effect of increases in the order processing delay from season to season, and the final interaction captures differences in this effect across Saturday 3pm matches (not televised) and non-Saturday 3pm matches (most matches televised). The 6 outcomes variables are as follows. Quotes is an indicator variable equalling 1 if there are quotes to both back and lay a bet in that second, Quoted Spread is defined in Section 3, Effective Spread is defined in Section 3, Depth is the available volume quoted at the best 3 back and lay prices, Order is an indicator variable equalling 1 if an order was executed in that second, and Volume is the trade size per second conditional on an order. Standard errors (in parentheses) are clustered at the bet-level, and ***, **, * and . indicates significance at the 0.1%, 1%, 5% and 10% level respectively.