Subsidizing Liquidity with Wider Ticks:
Evidence from the Tick Size Pilot Study

Robert P. Bartlett, III*
University of California, Berkeley

Justin McCrary**
University of California, Berkeley, NBER

Abstract:

Using data from the 2016-2018 tick size pilot study, we examine the efficacy of using wider tick sizes to subsidize market-making in small capitalization stocks. We demonstrate that realized spreads decay quickly within the initial microseconds of a trade, consistent with market makers being subject to adverse selection from asymmetrical speed among market makers. The effect reduces the subsidy offered by wider tick sizes, particularly for non-HFT market makers. The profit subsidy from wider tick sizes is also compromised by a significant shift in trading to "taker/maker" exchanges and to midpoint trading in non-exchange venues. The pilot's exception for midpoint trades also accounts for the fact that nearly a third of trading remains in non-exchange venues despite the inclusion of a trade-at rule. Overall, these findings point to considerable inefficiencies in the pilot study's goal of using wider tick sizes to subsidize liquidity provision in small capitalization stocks.

Draft Date: November 22, 2017
JEL codes: G10, G15, G18, G23, G28, K22
Keywords: tick size pilot; high-frequency trading; market making; market structure

***Statement of Financial Disclosure and Conflict of Interest*: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.**

* rbartlett@berkeley.edu, 890 Simon Hall, UC Berkeley, Berkeley CA 94720. Tel: 510-542-6646.
** jmccrary@berkeley.edu, 586 Simon Hall, UC Berkeley, Berkeley CA 94720. Tel: 510-643-6252.

I. Introduction

In October 2016, the Securities and Exchange Commission (SEC) commenced a two year pilot study that widens the minimum quoting increment—or "tick size"— from $0.01 to $0.05 for nearly 1,200 small capitalization stocks. Crafted in the wake of a steady decline in the number of listed companies on U.S. exchanges, the goal of the pilot study is to examine whether widening the tick size can enhance the liquidity and trading of small capitalization stocks.[1]

A central premise behind the tick size pilot is that "a widened tick increment could increase market maker profits and that the increased profits could foster a more robust secondary market for small capitalization stocks (and ultimately a more robust primary market) by, for example, increasing liquidity, enhancing the attractiveness of acting as a market maker, and possibly increasing the provision of sell-side research" (SEC 2015). In this regard, the pilot study is consistent with the theoretical literature on tick sizes that evaluates the regulatory choice of tick size as involving a tradeoff between minimizing transaction costs for investors and subsidizing liquidity providers to make a market in a security (Angel 1997). According to this theory, a wider tick size (such as existed prior to decimalization in 2001) should encourage dealers to make a market in a security because the market maker "pockets the spread" in a stationary market: the tick represents the minimum round-trip profit to a dealer who can buy at a lower bid price and sell at a higher offer price. Yet by the same reasoning, larger tick sizes increase transaction costs for liquidity takers who buy at the offer and sell at the bid. The tick size can thus be thought of as taxing liquidity takers in order to subsidize liquidity provision.

The notion that the market maker pockets the spread is a theoretical proposition, as it posits a stationary market. Empirically, evidence from the decimalization of stock prices in 2001

[1] The pilot study itself is a product of a mandate contained in the 2012 Jumpstart Our Businesses and Startups Act (JOBS Act) for the SEC to study the effects of the decimalization of stock prices in 2001 on IPOs and small and middle capitalization companies (Nallengara and Ramsay, 2013).

highlights how institutional features of the trading market can interfere with this straightforward relationship between tick sizes and market maker profits. For instance, while average quoted and effective spreads declined following decimalization (see, e.g., Chakravarty, Harris, and Wood (2001); Bacidore, Battalio, and Jennings (2003)), the effect was minimal for small capitalization stocks, which had large spreads prior to decimalization (Bessembinder, 2003). As another example, the smaller spreads of decimalization permitted dealers to offset any drop in profits through participating in more trades (Ronen and Weaver, 2001; Coughenour and Harris, 2004). These second order effects complicate understanding how tick sizes relate to liquidity providers' incentives.

In this paper, we use new data from the tick size pilot to examine the central premise in the theoretical literature on tick sizes and market making—namely, to what extent does widening spreads increase market maker profitability? We focus on this first order effect of wider tick sizes for the simple reason that market maker profitability is the primary channel through which wider ticks are believed to improve the liquidity of small capitalization stocks. Evidence that wider ticks result in greater market maker profitability is accordingly a necessary condition for wider ticks to achieve their stated policy objective. Conversely, evidence that market makers capture only a fraction of the enhanced transaction costs created by wider ticks would call into question the efficiency of using wider ticks to subsidize liquidity in small capitalization stocks.

We begin with an empirical finding. As noted by Rindi and Werner (2017), daily market maker profits reported to the SEC by FINRA have increased by approximately 40% in transactions involving securities that have been treated with nickel tick sizes—a notable increase in profitability. Yet, given that the tick size for these securities was quintupled, the increase in market maker profits is in many ways more modest than what proponents of the tick size pilot might have envisioned.

We posit that there are at least three reasons to question the efficacy of wider ticks in enhancing market maker profitability in contemporary equity markets. The first concerns the risk of informed trading in a context where liquidity provision is increasingly performed by high-frequency trading (HFT) firms. Consider, for instance, the framework of Aït-Sahalia and Saglam (2017) in which a strategic high frequency trader receives a signal about future order flows and exploits its speed advantage to optimize its quoting policy. For instance, if the bid-ask in a penny quoting environment is \$10.02 x \$10.08, a market maker receiving an incoming buy-order might fill the order at \$10.08 by selling short with an expectation of covering at the bid of \$10.02, earning 6 cents per share. In an Aït-Sahalia and Saglam (2017) framework, however, HFT liquidity providers observing the transaction may update their own buy orders rapidly, moving the best bid above \$10.02 before the market maker can cover. Such a result will force the initial market maker to quote at or better than this new price, reducing its expected profits.

The ability of HFT liquidity providers to adjust quotes in this fashion can jeopardize the market maker subsidy ostensibly created by wider spreads. For instance, if the original bid-ask spread were \$10.00 x \$10.10, our market maker's ability to capture the 10 cent spread will be compromised by the incentives of HFT liquidity providers to update bids to \$10.05 upon observing buying interest. Price-time priority rules observed by exchanges further complicate our market maker's ability to profit from the wider spreads: As HFT liquidity providers react to the buying interest, the wider ticks should create longer queue lines at the best bid of \$10.05, enhancing the challenge of covering the short position (Yao and Ye, 2014).

These long-queue lines point to a second reason to question the efficacy of wider ticks in enhancing market maker profitability. Using a regression discontinuity design, Bartlett & McCrary (2012) and Kwon et al. (2015) find that, relative to the \$0.0001 minimum price variation (MPV) for stocks that trade below \$1.00 per share, the penny MPV for orders priced above \$1.00 per share has the result of creating long queue lines at the national best bid or offer (NBBO). Because the MPV regulates quotes but not trades, however, traders seeking to avoid these long-queue lines can trade in non-exchange venues at sub-penny prices, generally by means of pegged mid-point orders (Bartlett & McCrary, forthcoming). In the prior example, for instance, our hypothetical liquidity provider—having gone short at \$10.10—may seek to cover its position by submitting a midpoint buy order to a dark pool rather than wait in line on a displayed venue at the best bid of \$10.05.

To the extent wider tick sizes result in longer queue lines, the migration of order flow away from exchanges towards midpoint orders in dark venues should accordingly place a further limitation on the effectiveness of wider spreads to subsidize liquidity providers, particularly those who display orders on exchanges. Indeed, concerns that wider tick sizes might drive trading to non-exchange venues induced the SEC to divide the pilot into three separate treatment groups. In group 1 ("TG1"), quotes must be priced in nickels, but trade prices are left unconstrained (similar to the current penny MPV rule.) In group 2 ("TG2"), quotes and trades must be priced in nickels. Finally, in group 3 ("TG3"), quotes and trades must be priced in nickels, and trading venues are also subject to a "trade-at" rule. The trade-at rule generally prevents price matching by trading centers that are not already displaying a quotation at that price. As such, the trade-at rule should have the effect of keeping trades on exchanges. However, as we document below, trades in both TG2 and TG3 are subject to various exceptions including one for midpoint orders, which potentially undermines the effectiveness of the trade-at rule given the common use of pegged midpoint orders in non-exchange venues (Bartlett & McCrary, forthcoming).

The final reason to question the efficacy of wider ticks in enhancing market maker profitability relates to exchange pricing models. Most exchanges compete for orders by using either a "maker/taker" or "taker/maker" price schedule. Under a maker/taker schedule, an exchange charges liquidity takers a per-share fee for removing liquidity from the exchange, a portion of which is used to pay a rebate to the trader providing liquidity. In contrast, a smaller group of exchanges use an inverse taker/maker schedule that pays the rebate to liquidity takers, while charging a per-share fee to the liquidity provider. As Angel, Harris, and Spatt (2015) document, inverted pricing schedules have the effect of creating a finer pricing grid than the minimum tick size. In effect, the liquidity provider is paying a liquidity taker to trade on the venue, reducing the net cost to the liquidity taker and increasing the probability a liquidity taker looks first to an inverted exchange. Thus, among exchanges, inverted exchanges should stand to gain market share in the pilot study for the same reasons that non-exchanges should gain market share. The fact that liquidity providers are assessed a fee for posting liquidity on inverted exchanges, however, also reduces the nominal subsidy offered by wider tick sizes.

We find that all three of these institutional features are empirically relevant in contemporary markets, which greatly compromises the pilot study's goal of subsidizing liquidity providers by using wider tick sizes. With respect to the role of informed trading, we document that in today's trading market, price impact begins to reduce realized spreads in the initial microseconds following a trade.[2] These results are consistent with wider ticks causing greater price impact due to HFT liquidity providers optimizing quotes in response to trades. Specifically, for securities

[2] As described below, we use realized spreads as our primary proxy for liquidity provider profits.

assigned to the control group in the pilot (i.e., penny priced quotes), our estimate of the price impact from a trade—measured by the percentage change in the NBBO midpoint following an observed trade—grows from approximately 2 basis points in the first microsecond after a trade to 10.7 basis points by the time one second elapses, a 4.6-fold increase.[3] However, we also show that for securities that are randomly assigned to the nickel treatment, price impact is greater and increases at a faster rate in the microseconds after a trade. For instance, for securities assigned to TG1, our estimate of price impact grows from approximately 2.8 basis points in the first microsecond after a trade to 20.4 basis points by the time one second elapses, an increase of more than 700%. Moreover, approximately 60% of this change occurs within the first millisecond of a trade. In contrast, approximately 52% of the one-second price impact change occurs in the first millisecond for the control group. Overall, these findings highlight how the HFT quoting environment can undermine efforts to subsidize liquidity providers using wider tick sizes. They also underscore the fact that to capture the wider spread, liquidity providers need to be faster than human market makers are capable of being, highlighting that any movement to a nickel quoting environment is likely to favor HFT market making.

We additionally document the extent to which the longer queue-lines associated with a nickel tick size can interact with the competition for order flow among exchange and non-exchange trading venues. Among exchanges, we find that inverted exchanges have experienced a significant increase in quoting activity and trading volume relative to maker/taker exchanges. Likewise, consistent with our predictions, we document notable changes in both the frequency and the form of non-exchange trading. Within TG1 and TG2, for instance, the size-weighted

[3] When a trade fills an existing quote on an exchange, the exchange produces a trade report and a quote report stamped to the same microsecond, reflecting the trade and the change in the exchange's best bid or offer, respectively. As such, we attribute the price impact at one microsecond to trades absorbing the displayed depth at the NBBO, forcing a mechanical change to the NBBO. We discuss this phenomenon in more detail in Section 6.

fraction of trades occurring in non-exchange venues increases from approximately 42% to approximately 46% and 44%, respectively. The result is the opposite in TG3 where non-exchange trading drops due to the trade-at rule, yet it still accounts for 32% of all trades.

Finally, we find that the nickel tick size results in a significant increase in the incidence of midpoint trades in both exchange and non-exchange venues. For instance, among non-exchange venues, the incidence among midpoint trades increases by 14% and 22.5% for securities in TG1 and TG2, respectively. Because of the trade-at rule, the effect is even more pronounced for securities within TG3, where the incidence of midpoint trades grows to over 65% of all non-exchange trades—an increase of nearly 200%. As such, the existence of the midpoint exception provides an important explanation for why nearly one-third of trades in TG3 remain in non-exchange venues despite the trade-at rule. For the same reason, it greatly limits the effectiveness of the nickel tick size to deliver a subsidy to providers of displayed liquidity.

Our findings have several implications for researchers and policy-makers. For researchers, our results provide compelling evidence that the institutional features of trading markets can interact with the tick size in ways that substantially reduce the ability of wider ticks to enhance liquidity provider profits. Moreover, by examining realized spreads beginning with the first microsecond after a trade, our study complements recent work by Conrad and Wahal (2017) who examine the term structure of realized spreads from 100 milliseconds to 600 seconds after a trade. They similarly emphasize that in electronic markets, realized spreads should be measured with much shorter time frames than has been the conventional practice. Using recent changes to the timestamps of the publicly available TAQ data, we show both the feasibility and the empirical relevance of measuring realized spreads from the first microsecond following a reported trade. This empirical strategy is also consistent with the theoretical work of Aït-Sahalia and Saglam (2017) who demonstrate that market makers are subject to adverse selection risk arising from both information asymmetries and speed asymmetries.

Finally, for policy-makers, our findings point to an inherent inefficiency in trying to use the tick size to subsidize liquidity provision for small company securities. Whatever merit this approach may have had in the past, the speed with which contemporary prices react to trades makes it extraordinarily difficult for all but the fastest traders to capture the subsidy created by the nickel tick size. To the extent the tick size pilot was motivated by a desire to reverse the effects of the decimalization of stock prices in 2001 (see, e.g., Weild, Kim and Newport, 2012), the policy accordingly suffers from a fundamental anachronism. Moreover, while liquidity takers have reduced some of the costs of wider ticks through exploiting the midpoint exception, they nevertheless pay effective spreads that are on average more than twice what they were in the penny environment. Thus, while market maker profits may have increased, our findings make it questionable whether the pilot study provides sufficient evidence that widening tick sizes would pass the type of cost-benefit analysis increasingly expected of financial reforms.

This paper proceeds as follows. Section 2 situates our study in light of other papers examining the tick size pilot. Section 3 provides an overview of our research design, describing both the tick size pilot and our use of the new microsecond timestamps. Section 4 discusses our sample of tick size securities and provides, to our knowledge, the first formal test of whether the tick size pilot study produced a balanced set of control and treatment securities. Section 5 discusses our estimation strategy, and Section 6 presents our results. Section 7 concludes.

II. Prior Studies of the Tick Size Pilot

Several papers have begun to examine the results of the tick size pilot. To our knowledge, we are the first to exploit the new microsecond timestamps in the TAQ data to illustrate why today's trading institutions work to undermine the capacity of wider ticks to subsidize liquidity providers. Griffith and Roseman (2017) and Rindi and Werner (2017) both examine how the pilot affects liquidity provider profits by examining realized spreads in pilot securities. Griffith and Roseman (2017) estimate that the aggregate costs of liquidity demanders increased by over $9,000 per stock day in test group securities with pre-pilot quoted spreads that are less than $0.05, while liquidity supplier profits increased by over $3,700 per stock day. However, in estimating realized spreads, they utilize a 5-minute post-trade window. We demonstrate why these estimated profits are so low by the fifth minute following a trade and, more importantly, why the majority of the subsidy available from wider ticks begins to decay within the initial microseconds after a trade.

Rindi and Werner (2017) likewise find that spreads (quoted and effective) widen after the tick size increase, as do realized spreads for all but the securities assigned to TG3. Similar to our approach, they also examine the term structure of realized spreads and estimate realized spreads at 30, 60, 120, 180, 240, and 300 seconds after a trade. However, as with Griffith and Roseman (2007), by beginning at 30 seconds after a trade, they do not observe how price impact disproportionately grows among all three treatment groups within the initial microseconds of a trade, and thus, cannot observe how the tick size pilot is affected by the adverse selection risk posed by high speed market makers. We also find that realized spreads increase for all three treatment groups, including TG3.

Three other papers expressly examine how the pilot has affected the incidence of market activity across competing venues. Comerton-Forde, Gregoire, and Zhong (2017) and Lin (2017) find that the first two treatment arms have increased trading on non-exchange venues, while TG3 drives trading to exchanges. Among exchanges, these papers also show that exchanges using inverted pricing models gain in market activity at the expense of maker-taker venues. Cox (2017) focuses exclusively on migration among exchange trades and similarly finds that the pilot has driven trading from maker/taker exchanges to inverted exchanges.

None of these papers examine the relationship of these observed effects within the context of the pilot's goal of increasing market maker profitability. Moreover, none of these papers examine the prediction made by Bartlett & McCrary (forthcoming) in the context of the change in tick size at the $1.00 order price that the pilot study would increase the incentive to trade at the midpoint of the NBBO. By expressly grappling with the midpoint exception, we help explain why, despite a quintupling of the tick size, average daily market maker profits increase by just 40%. Moreover, Comerton et al. (2017), Lin et al. (2017), and Cox et al. (2017) each examine trading only with respect to the final months of 2016. These sample selection choices may explain why Lin (2017) finds a decrease in trading across maker/taker venues, while Comerton (2017) and Cox (2017) find no decrease in maker/taker trading volume. Our larger sample period and methodological approach reveals an increase in trading volume among inverted exchanges and a decline among maker/taker exchanges across all three treatment groups.

III. Overview of Research Design

To assess empirically the extent to which a nickel environment subsidizes market markers in small capitalization firms, we exploit two recent policy interventions. The first policy intervention is the tick size pilot itself, which utilizes a randomized experimental design to explore the effects of imposing a nickel tick size on small capitalization stocks. The second is the introduction of new microsecond timestamps in August 2015 by the two Securities Information Processors (SIPs), which permit us to estimate market makers' profits in the microseconds following an observed trade.[4] We provide institutional details regarding both interventions in this section.

A. Tick Size Pilot

In all of our analyses, we rely on the fact that in implementing the two-year tick size pilot study, the SEC utilized a randomized controlled trial to study the effects of widening the tick size from a penny to a nickel. Under the terms of the pilot study, the SEC used a measurement period (the "Measurement Period") commencing three months prior to September 2, 2016 to determine which exchange-listed securities were eligible for inclusion either as a control security or a security subject to a nickel tick (collectively, the "Pilot Securities"). Specifically, a security was eligible for inclusion as a Pilot Security if (a) on the last day of the Measurement Period, the security had an aggregate market capitalization of $3 billion and a closing price of at least $2.00 per share, and (b) during the Measurement Period, the security had a closing price on every U.S venue of at least $1.50 per share, a consolidated daily average volume of 1 million shares or less, and a volume-weighted average price of at least $2.00 per share. In total, 2,399 securities were identified as Pilot Securities.[5]

Following the identification of Pilot Securities, an Operating Committee consisting of representatives from all exchanges and FINRA assigned each Pilot Security to either a control group or one of three treatment groups. As noted previously, the SEC utilized three different treatment groups in light of the concerns noted previously that widening the tick size might induce traders to trade at subpenny prices in non-exchange venues given that the tick size rule

---

[4] All exchanges and brokers are required to submit quote updates and trade reports occurring on a venue to one of two SIPs. Quotes and trades in NYSE-listed securities ("Tape A" securities) and securities listed on regional exchanges and their successors ("Tape B" securities) must be submitted to the Securities Industry Automation Corporation ("SIAC"), a subsidiary of the NYSE which acts as the central SIP for any transaction in Tape A and Tape B securities. Quote updates and trade reports in Nasdaq securities must be sent to the SIP managed by Nasdaq. The quote updates and trade reports sent to these two SIPs constitute the publicly available TAQ data.

[5] The original list of Pilot Securities can be found at: http://www.finra.org/sites/default/files/Tick_Size_Pilot_Selection_Process.pdf.

conventionally applies only to quotations and not to trades. As such, Pilot Securities were assigned to one of the following four groups:

- **Control Group**: Securities in this group are quoted at their current tick size increment of a penny and remain subject to the existing rule that trades can be made in penny and subpenny prices.
- **Treatment Group 1 (TG1)**: Securities in this group are quoted in \$0.05 increments (subject to limited exceptions), but otherwise can be traded in pennies or sub-pennies.
- **Treatment Group 2 (TG2)**: Securities in this group are quoted in \$0.05 increments (subject to limited exceptions) and must also trade in \$0.05 increments (subject to exceptions, including executions at the midpoint of the NBBO, certain retail investor executions and negotiated trades.)
- **Treatment Group 3 (TG3)**: Securities in this group are subject to the quoting and trading requirements of TG2 and are additionally subject to a "trade-at" requirement, which generally prevents price matching by trading centers that are not already displaying a quotation at that price (subject to certain exception, including those for TG2.)

Speaking loosely, TG1 can be thought of as "quotes in nickels," TG2 can be thought of as the "quotes and trades in nickels," and TG3 can be thought of as the "quotes and trades in nickels, and the trade-at rule."

Operationally, in assigning securities to one of these groups, the Operating Committee first assigned all Pilot Securities to a stratum, after which securities were selected at random within stratum for inclusion in one of the three treatment groups. The strata were defined based on a security having: (1) a low, medium, or high share price based on the value-weighted average price during the Measurement Period, (2) a low, medium, or high market capitalization based on the last day of the Measurement Period, and (3) a low, medium, or high trading volume based on the average daily trading volume during the Measurement Period. In each of these three classifications, Pilot Securities were classified based on the tercile (i.e., low, medium, or high) in which it fell.

In all, this classification scheme resulted in Pilot Securities being assigned to a total of 27 strata of Pilot Securities (i.e., H-H-H, H-H-M, etc.). Because some of these strata had a small number of securities, however, the Operating Committee elected prior to randomization to combine some of the smaller strata, resulting in 21 "revised strata." Randomization yielded approximately 400 securities in each of TG1, TG2, and TG3. Those Pilot Securities not placed into the three treatment groups constituted the control group.

Announcement of the Pilot Securities and their assignments was publicly made on September 6, 2016. At the time of the initial assignment, 397 (16.66%) securities were assigned to TG1, 395 (16.58%) securities were assigned to TG2 and TG3, and 1,196 (50.19%) securities were assigned to the control group.[6] The tick size pilot commenced on October 3, 2017 as securities were phased into their assignment treatment over the ensuing weeks with full implementation occurring on October 31, 2016.[7]

B. Microsecond Timestamps

Because we seek to examine the quoting environment at the time of a trade as well as in the microseconds following it, our design also exploits the fact that since August 2015, the TAQ data records in microseconds the precise time at which a quote update or trade report occurred on a trading venue (see Bartlett & McCrary, 2017 for a discussion). These new data permit a novel look at how adverse selection costs affect liquidity providers in today's trading environment.

In the classical setting, market makers face the possibility that a liquidity taker is better informed about the fundamental value of the security, generating adverse selection risk for the market maker (see Glosten and Milgrom, 1985). For instance, a market maker might sell to a buyer who knows that posted asks undervalue the security, resulting in an overall increase in the

[6] See http://www.finra.org/industry/test-group-assignments. While we have been unable to confirm the probability of selection and in particular whether it was constant across strata—the key condition which would obligate any analysis to condition appropriately on strata—there is little evidence of meaningful departures from a constant probability of selection. However, conditioning on strata here guards against bias and substantially improves the precision of estimated treatment effects. Consequently, as discussed in more detail below, we fully condition on strata in all of our analyses.

[7] The rollout of the plan during October 2016 is described more fully at the following link: http://www.finra.org/industry/tick-size-pilot-program-implementation-plan.

NBBO following the trade as the informed buyer picks off underpriced orders. Such a market maker will then have to cover the short position by buying at the higher national best bid, potentially causing a loss for the market maker. Given this adverse selection risk, conventional approaches to examining the effect of adverse selection costs on liquidity provider profits have focused on decomposing effective spreads into a *realized spread* at time horizon *t* that is captured by a liquidity provider and a *price impact* at time horizon *t* that reflects the post-trade movement in the quote midpoint that undermines a liquidity provider's ability to capture the full spread.

Our framework differs from this classical setting insofar that we follow Aït-Sahalia and Saglam (2017) in assuming that price impact can also arise from HFT liquidity providers who rapidly update quotes in response trades. To illustrate, consider again the example provided in the Introduction in which a stock was quoted at $10.02 x $10.08, during which a market maker filled an incoming buy-order by selling short a share at $10.08. Assuming a post-trade movement of the NBBO to $10.04 x $10.08, the market maker that covered its short position at the new midpoint of $10.06 would earn 2 cents (its realized spread) rather than 3 cents (the realized spread had she covered at the prior midpoint of $10.05). The 1 cent increase in the quote midpoint reflects the price impact associated with this hypothetical trade.

Estimating price impact from this form of adverse selection naturally raises an empirical challenge given the need to examine price impact in the microseconds following a trade. Prior to the introduction of the SIPs' new microsecond timestamps, researchers using TAQ data to estimate price impact and realized spreads faced the practical challenge that the timestamp assigned to a quote or trade report was made in milliseconds and reflected the time that one of the two SIPs finished processing the transaction report received from the reporting venue. Given reporting latencies in the quote and trade reports sent to the SIPs, researchers were accordingly required to make assumptions about the quoting environment at the time of a trade to estimate effective spreads and price impact (see Bessembinder 2003 for a discussion). These timestamp issues also posed challenges for understanding the quoting environment in the moments after a trade. However, conventional approaches to measuring realized spreads (including the formal definition of realized spreads in SEC Rule 605) generally averted these challenges by using a 5-minute lag (or longer) following a trade for measuring realized spreads and price impact (Bessembinder 2003; Goyenko et al. 2009).

In contrast, by using the new microsecond timestamps, we can empirically estimate the price impact and realized spreads following a trade from the initial microsecond after it occurs through any point in time. As such, our approach permits not only an evaluation of whether a wider tick size produces enhanced adverse selection due to HFT optimizing, but also whether realized spreads today are more appropriately measured in the microseconds and seconds following a trade, rather than in the minutes following it.

IV. Data

A. Sample Construction

We obtain our sample data from the trade and quote reports from the NYSE TAQ data for Pilot Securities from March 7, 2016 through June 30, 2017. Due to the implementation of the pilot study, we collect data across the following four time periods established by the pilot:

March 7, 2016 – September 2, 2016:	Pre-Pilot Period
September 6, 2016 – September 30, 2016:	Pilot Securities Announcement Period
October 3, 2016- October 31, 2016:	Pilot Phase-In Period
November 1, 2016-June 30, 2017:	Full implementation

Pre-Pilot Period. We commence with March 7, 2016 as this is the first date on which FINRA published a list of pre-pilot securities that were to be used to evaluate the pre-pilot trading environment. Under the plan, trading centers and market makers were required to collect trading data for these pre-pilot securities until September 2, 2016, after which the data collection requirements would apply only to the Pilot Securities. As discussed below, we use measures collected during the Pre-Pilot Period as a control in some of our analyses.

Pilot Securities Announcement Period. On September 2, 2016, the identity of Pilot Securities was announced by Nasdaq and the NYSE following a three-month measurement period used to identify exchange-listed securities that were eligible to be Pilot Securities. Assignment of Pilot Securities to control and treatment groups was published by Nasdaq and NYSE on September 6, 2016.

Pilot Phase-In Period. Implementation of the pilot occurred on five different Mondays during October 2016. Securities from TG1 and TG2 were phased in on October 3, October 10, and October 17. Securities from TG3 were phased in on October 17, October 24, and October 31.

Full Implementation. Under the plan, all securities assigned to a treatment group were subject to the experimental treatment as of November 1, 2016. Data in our sample covering this time period accordingly run from November 1, 2016 through June, 2017.[8]

Given the technical challenges of implementing the pilot study, we begin our sample period on November 7, 2016, giving market participants a one week transition period from the date of full implementation. Our sample period ends on June 30, 2017. In all analyses, we focus on those Pilot Securities that traded for the duration of this sample period. Due to acquisitions, delistings, and other corporate actions following the commencement of the pilot study, the total number of Pilot Securities that satisfies this criteria is 2,168 securities.

[8] The final day of the pilot will be October 3, 2018, followed by a post-pilot data collection period that runs until April 1, 2019.

Following Rindi and Werner (2017), our sample selection criteria also reflects the fact that the treatment effect—and therefore, the possibility for a market-maker subsidy—will generally be concentrated in those securities where quoted spreads tend to be less than 5 cents. Yet the selection criteria for Pilot Securities implicitly focus on less liquid stocks, suggesting that many of these securities should be expected to trade at spreads in excess of 5 cents. In Table 1, we examine average quoted spreads for Pilot Securities across control and treatment groups during the four time periods noted previously, classifying securities into whether they had an average pre-pilot quoted spread of less than 5 cents. In our analyses below, we limit our sample to those Pilot Securities having average pre-pilot quoted spreads of less than 5 cents in light of the fact that any rents provided to market makers by virtue of widening the tick size from a penny to a nickel will be largely concentrated in these securities.

[Insert Table 1]

In total, these restrictions result in a core sample of 1,068 securities, consisting of 535 securities assigned the control group (50.09%), 175 securities assigned to TG1 (16.39%), 170 securities assigned to TG2 (15.92%), and 188 securities assigned to TG3 (17.60%). These proportions are generally in line with the original assignment of Pilot Securities discussed in Section III(A).

Additionally, to ensure that all quotes and trades occur during the trading day after the opening cross and before the closing auction, we subset the data to exclude quotes and trades occurring before 9:45:00 and after 15:44:59.999999. As noted in Holden and Jacobsen (2014), the NBBO file of the Daily TAQ file is incomplete; therefore, we manually calculate the NBBO for each security for each microsecond during our sample period using quote updates from the daily TAQ data and the standard Hasbrouck algorithm. In so doing, we restrict our analysis to those quotations that are eligible to establish an exchanges' best offer or best bid (i.e., quotation updates having a condition of A, B, H, O, R, W, or Y).

B. Balance Tests and Summary Statistics

In Table 2, we present summary statistics for our sample of 1,068 Pilot Securities using a variety of proxies for a security's liquidity during the Pre-Pilot Period, dividing our analysis separately for control and treatment arms. We do so to ensure that our sample selection criteria produces a balanced set of control and treatment securities. There is a natural question regarding how precisely we estimate treatment-control differences given that the randomization was not simple and instead depended on stratification. We estimate differences between TG1, TG2, TG3, and control in a manner consistent with how we approach estimating treatment effects below. As we discuss there, the method we use is the most efficient (in the statistical sense of the term) method of estimating treatment effects that imposes only the knowledge that treatment was assigned randomly conditional on strata, and this implies that the balancing tests we report are the most powerful tests possible without adopting parametric assumptions.

[Insert Table 2]

The results of these balancing tests are shown in Table 2 where we examine fifteen different market measures across control and treatment groups. Estimates for each measure are based on observed trades from the Pre-Pilot Period and are obtained using the same regression framework discussed below in Section 5. As such, the control group mean reflects a regression-adjusted size-weighted average for securities assigned to the control group, while estimates for each treatment group reflect the estimated treatment effect for that group. As Table 2 shows, the difference between the treatment and control assignment is statistically insignificant for 14 of the 15 measures. The single exception is for the mean fraction of trades marked as being part of an

Intermarket Sweep Order, which was slightly lower for each treatment group during the Pre-Pilot Period. The effect, however, was only moderately insignificant. Given the large number of tests showing no significant difference between control and treated securities, and the prior knowledge regarding the protocol used by the Operating Committee, we conclude that Table 2 is not inconsistent with there having been true random assignment conditional on strata. This assumption thus forms the basis for our estimation of the effects of the tick size pilot below.

C. Compliance with Treatment Assignment

Our final preliminary test examines whether market participants complied with the treatment assignments of the Pilot Study. We begin by plotting in Figure 1A the incidence of nickel-priced NBBO updates over time by treatment assignment, which reveals a stark increase in nickel-priced updates as Pilot Securities were phased into treatment groups. The three vertical lines in the figure indicate the dates corresponding to the seams between the four different time periods associated with the pilot: September 2, 2016, September 30, 2016, and one week after Halloween 2016 (November 7, 2016), when our sample period begins. Overall, the evidence is consistent with participants complying with the nickel-quote rule applicable to securities assigned to TG1, TG2, and TG3. In Figure 1B, we present a similar plot showing the incidence of nickel-priced trades, which similarly reveals a large increase in nickel priced trades across all the three treatment groups. Somewhat surprisingly, Figure 1B also reveals a slight increase in nickel priced trades among control securities immediately after full implementation, although the effect diminishes after the initial weeks of the Pilot Study.

Overall, these tests confirm that market participants complied with the treatment assignments as required by the Pilot Study.

[Insert Figures 1A and 1B]

V. Estimation Strategy

Given its experimental design, the pilot study permits a unique opportunity to identify the effects of the treatment assignments on trading activity. However, the structure of the pilot also poses a number of methodological questions—a point that has yet to be examined in prior work examining the Tick Size Pilot. For example, while the Tick Size Pilot Program Implementation Plan[9] provides extensive description of the pilot as having involved randomization, a careful examination of the record reveals that randomization was conditional, as opposed to unconditional. This means that a simple comparison of average outcomes between the treatment and control groups does not necessarily identify the average treatment effect (ATE) typically delivered by an unconditionally randomized evaluation. Instead, attention must be paid to the conditioning set upon which the Operating Committee relied before randomly selecting stock from within strata.

When there are departures from pure randomization, in addition to issues of identification, there are also issues of efficiency to be considered. Traditional econometric approaches to conditional randomization would assume constant coefficients and would approach the problem using regression with indicators for treatment arms as well as strata indicators. However, the literature on treatment effects has established that greater efficiency can be obtained from somewhat different estimators (Hahn 1998, Imbens 2004). In the context of randomization conditional on strata, there are equivalences among the three leading approaches to estimation (regression, matching, and reweighting). The semiparametric efficient estimator that is easiest to explain is matching: within each stratum, we compute differences in mean outcomes between treatment arms. Those stratum-specific differences are then averaged using the number of stocks

[9] See http://www.finra.org/industry/tick-size-pilot-program-implementation-plan.

in each stratum as a weight.[10] The resulting estimator identifies the ATE and is additionally semiparametrically efficient.[11]

Nonetheless, obtaining valid standard errors for this estimator can be challenging computationally. Following Wooldridge (2010, Chapter 21), we use a computational trick that allows for recovery of the efficient semiparametric estimator from a simple regression. This involves de-meaning the set of (all but one) strata indicators prior to forming interactions between the three treatment arm indicators with those strata indicators. A regression of the outcome on treatment arm indicators, strata indicators, and the interaction terms then yields coefficients on the treatment arm indicators that are numerically equivalent to the matching estimator described above. Because this particular regression estimator is numerically equivalent, it also must be semiparametrically efficient.

In addition, the constant term from that regression recovers the efficient estimator of the outcome under the control regime. This latter step is practically useful, since we consider not just linear effects in which we compare the difference in the expectation of counterfactual outcomes, but also nonlinear effects such as the ratio of the expectation of counterfactual outcomes.

Finally, we note that while our regressions are based on large scale data sets, such as quotes or trades, the regressions involve a covariate grouping structure: the covariates only take on a fixed number of values, since neither stratum nor treatment assignment varies except cross-sectionally across stocks. Consequently, we focus on cross-sectional regressions where the

[10] Somewhat more formally, suppose that the strata are indexed by j with average outcomes under treatment group t within stratum j given by $Y_j(t)$ and number of stocks in the stratum of n_j. Then the matching estimator of the average treatment effect t versus t is the weighted average across the 21 strata of $Y_j(t)$- $Y_j(t)$ using weights n_j.

[11] Note that we ignore compliance problems here and adopt a perspective that we seek to estimate an intention-to-treat (ITT) parameter. This is largely because compliance is high enough that we do not believe there will be a meaningful discrepancy between an ITT parameter and a ATE parameter or, for that matter, an ITT parameter and a local average treatment effect (LATE) parameter.

outcome variable is an average of the underlying data, weighted as appropriate (e.g., by size or value). In those cross-sectional regressions, we generally give each stock in the sample equal weight.

VI. Results

In this section, we analyze the effectiveness of using the tick size pilot to enhance liquidity provider profits. As noted previously, a central goal of the pilot was to increase liquidity provider profits through increasing the quoting increment from a penny to a nickel, which was anticipated to lead to wider quoted and effective spreads. The wider tick size was also anticipated to increase quoted depth at the NBBO given that traders would have fewer price points at which to express trading interest. Quoted depth was also anticipated to increase for TG3 on account of the trade-at rule. Accordingly, we begin with an analysis of the extent to which the pilot produced these anticipated effects.

A. Quoted Spreads, Effective Spreads, and Quoted Depth

In Panel A of Table 3, we present the results for quoted and effective spreads. The first two columns provide results for quoted spreads. For each security, we estimate spreads at the microsecond of a trade and present the size-weighted mean for the control group and each treatment group. Column one presents the estimates over our sample period for the control group, using the regression framework discussed in Section V. For each treatment arm, b_1-b_0 provides the gross change in average spreads over the control group mean, and b_1/b_0 estimates the percentage increase in average spreads relative to the control group mean. As expected, quoted spreads increased considerably across all three treatment groups. Compared to the control group mean of approximately $0.03 per share, quoted spreads for TG1, TG2, and TG3 increased by $0.0291, $0.0292, and $0.0270 per share, respectively. For securities in our sample, the pilot thus had the effect of virtually doubling the average quoted spread.

[Insert Table 3]

To obtain additional precision of our estimates, we also estimate quoted spreads including as an additional control a security's average quoted spread during the Pre-Pilot Period. As noted, our regressions are generally cross-sectional regressions.[12] These results, presented in the second column, are largely the same. However, as suggested by the 0.7019 coefficient for these pre-period effects, a security's quoted spread during the Pre-Pilot Period was strongly predictive of its quoted spread during our sample period. Adding a control for these pre-pilot outcomes also provides a notable improvement in the regression fit, as reflected in the increase in R^2 from 0.66 to 0.75. In light of the additional precision obtained by using these pre-period controls, we focus our discussion on our results that utilize pre-period controls for the outcome variable of interest.

In columns 3 and 4, we present results for relative quoted spreads, which we calculate as the difference between the natural log of the NBO and the natural log of the NBB as of an observed trade. Calculating quoted spreads in this fashion is accordingly a first order approximation of calculating the quoted spread relative to the NBBO midpoint. As with gross quoted spreads, the relative quoted spreads increased considerably across all three treatment groups during the sample period. Specifically, while the control group had a mean quoted spread that was 0.3101% of the NBBO midpoint, TG1 had a mean quoted spread that was 39.79 basis points higher, while TG2 and TG3 had mean quoted spreads that were higher by 36.2 basis points and 41.24 basis points, respectively. Relative to the control group, this translates into an increase in relative quoted spreads of 128% for TG1, 117% for TG2, and 133% for TG3. Focusing on the

[12] To accommodate the pre-period quoted spreads, we simply computed the average quoted spread corresponding to the Pre-Pilot Period. That yields an additional cross-sectional covariate to be added to our cross-sectional regression approach.

difference between the treatment and control (b_1-b_0) for each treatment group, the size of the standard errors indicates that each treatment group was significantly different from the control group mean (p<0.01).

In columns 5 and 6, we present results for relative effective spreads. We calculate relative effective spreads for buy orders by multiplying by two the difference between the log of an observed trade price and the log of the NBBO midpoint prevailing at the moment of the trade; for sell orders, we take two-times the difference between the log of the NBBO midpoint and the log of the trade price. We classify trades as having been buy- or sell-side initiated using the Lee and Ready (1991) algorithm with no lag (see Bessembinder and Venkataraman, 2010). As with relative quoted spreads, all three treatment groups witnessed a sizeable increase in relative effective spreads. In particular, TG1 had a mean relative effective spread that was 31.69 basis points higher than the control group mean, while TG2 and TG3 had mean relative effective spreads that were higher by 28.50 basis points and 27.77 basis points, respectively. Relative to the control group, this translates into an increase in relative effective spreads of 128% for TG1, 115% for TG2, and 112% for TG3. As with quoted spreads, the size of the standard errors indicates that each treatment group was significantly different from the control group mean.

A notable difference between quoted and effective spreads concerns the rank ordering of the estimated treatment effects across the three treatment arms. For instance, while TG3 had the highest relative quoted spreads, it had the lowest relative effective spreads. In unreported tests, we formally test the hypothesis of equal treatment effect across all three treatment arms, as well as equal treatment between any two treatment groups. In none of these tests can we reject the hypothesis of equal treatment. Nonetheless, we note here the differential rank ordering, as it highlights the importance of evaluating the tick size pilot in light of how traders respond to wider ticks by searching for price improvement in non-exchange venues. Somewhat surprisingly, Table 3 suggests this phenomenon is most dramatic for TG3, notwithstanding the fact that the trade-at rule should result in the greatest share of trades occurring on exchanges at nickel prices. We explore this issue in more detail in Section 6(C).

Finally, in Panel B we provide estimated treatment effects for quoted depth at the NBBO. We likewise measure inside quoted depth as of an observed trade, and we define it as the natural log of one-half the sum of quoted inside depth at the bid and quoted inside depth at the ask. As expected, all three treatment arms experienced a significant increase in quoted inside depth, which is consistent with the fact that the nickel quoting increment forces traders to express trading interest across a fewer number of price points. The increase is especially pronounced for TG3, where average inside depth increased over 400% relative to the control group mean. This latter effect is consistent with the expected effects of the trade-at rule given that the rule generally prevents price matching by trading centers that are not already displaying a quotation at that price. Traders seeking to trade in a TG3 security therefore have incentives to display quotes at the NBBO, increasing quoted depth. As we discuss below, however, the increase in quoted depth can create strong incentives to utilize the midpoint trade exception to the trade-at rule, potentially explaining why effective spreads for TG3 are no higher than those of TG1 and TG2.

B. Realized Spreads

While the prior section reveals that relative quoted spreads and relative effective spreads more than doubled for all three treatment groups, it would be a mistake to conclude that these are the spreads liquidity providers expect to earn by virtue of a nickel quoting rule. As noted previously, a primary reason relates to the likelihood that a security's price will move against a liquidity provider as she seeks to capture the spread. The extent of this post-trade movement in the price of a security—or a trade's price impact—effectively reduces the portion of the trade's effective spread that can be captured by a liquidity provider.

In keeping with the literature, we therefore estimate profits to liquidity takers from the tick size pilot by decomposing effective spreads into a realized spread component and a price impact component. Formally, we define the *i*th trade's realized spread at time horizon *t* as

$$Realized\ Spread_i(t) = 2D_i \text{ x } (\ln P_i - \ln N_i(t))$$

where P_i represents the *i*th trade's observed execution price, $N_i(t)$ represents the midpoint of the NBBO for the security at *t* (potentially fractional) seconds after the time of the *i*th trade, and D_i represents a trade's direction (buy/sell) using Lee-Ready. Conversely, we define a trade's price impact at time horizon *t* as

$$Price\ Impact_i(t) = 2D_i \text{ x } (\ln N_i(t) - \ln N_i(0))$$

where $N_i(0)$ represents the midpoint of the NBBO at the precise moment of the *i*th trade. As with our definition of relative effective spreads, our use of log-differences has the effect of measuring both realized spreads and price impact on price-adjusted basis. Specifically, for realized spreads, the measure produces a first order approximation of the realized spread for a trade as a percentage of the observed price of the trade; for price impact, the measure produces a first order approximation of the price impact for a trade as a percentage of the quote midpoint in effect at the time of an observed trade. As we discuss in Section 3(B), we exploit the new microsecond timestamps to estimate these measures commencing with the first microsecond following a trade (*t*=0.000001), and for every microsecond thereafter through the fifth minute following an observed trade (*t*=300).

In Figure 2A, we present a scatter plot of average price impact for each treatment group as well as for the control group as a function the natural log of *t*, the time that has elapsed since an observed trade. Figure 2B presents a similar scatterplot illustrating average realized spreads as a function of the natural log of *t*. As one might expect, the two figures are almost exactly the opposite from one another with respect to the time-series behavior they present. For instance, in the first microsecond following a trade, average price impact is near zero across all three treatment arms, while average realized spread ranges from approximately 23 to 26 basis points depending on the treatment arm. The results reflect the fact that in the first microsecond after a trade, market participants are, on average, unlikely to have had sufficient time to respond to the trade to move prices, diminishing the risk of price impact. Moreover, because quoted prices are unlikely to have changed, a trader who provided liquidity for the trade retains the opportunity of capturing nearly all of the effective spread by completing an off-setting passive trade at the NBBO.

At the same time, while price impact is nearly non-existent in the first microsecond following a trade, it is worth emphasizing that it nevertheless exists even within this time span following a trade. While it is possible this result reflects the presence of informed traders, we suspect the primary explanation stems from trade and quote reporting conventions among exchanges. In particular, when a trade fills an existing quote on an exchange, the exchange matching engine produces separate trade and quote reports that are each stamped to the same microsecond to reflect the trade and the change in the exchange's best bid or offer, respectively. If an exchange happens to hold the remaining depth at the NBBO at the time of a trade, the transaction will accordingly produce a "mechanical" change in the NBBO by the first microsecond of the trade.

Notably, the existence of this mechanical effect is itself informative of how the wider tick size is affecting the liquidity of pilot securities. Despite the considerably greater quoted depth for all three treatment groups, this mechanical effect is higher for treated securities than for control securities. This suggests that, despite the greater inside depth, trades in treated securities nevertheless dislodge the NBBO. Moreover, conditional on the NBBO being dislodged in this fashion, the mechanical effect will be larger when the NBBO is quoted in nickels rather than pennies given that the second level of prices will differ by a nickel.[13]

Figures 2A and 2B also reveal that the likelihood a liquidity provider can capture the wider spreads caused by the tick size pilot begins to dissipate rapidly in the moments following a trade. The effect is especially strong with respect to trades subject to the trade-at rule in TG3. By the first second following a trade, Figure 2B indicates that average realized spreads are nearly one-half of their initial size for this group, while realized spreads for TG1 and TG2 are roughly two-thirds of their initial size. Thereafter, our estimate of average realized spreads across all three groups continues to decline, though at a more modest rate through the fifth minute following a trade.

In Table 4, we further explore the decay rate of realized spreads by examining our estimate of realized spreads at three distinct points in time following a trade: at the first microsecond (t=0.000001), at the first second (t=1) and at the fifth minute (t=300). In column 1, we first present the value-weighted average relative effective signed spread for the control group and for each treatment group. A trade's relative effective signed spread is equivalent to one-half of the

[13] In unreported tests, we evaluate the fraction of trades in our sample in which a buy (sell) transaction results in a higher (lower) NBBO in the first microsecond following a trade. Overall, approximately 13% of trades within the control group produce such a change in the NBBO, with the average amount of the change being \$0.016 higher (lower). In contrast, approximately 8% of trades in one of the three treatment groups produce such a change in the NBBO, with the average amount of the change being \$0.05 higher (lower). Accordingly, while the greater depth across the three treatment groups means that fewer trades dislodge the NBBO, the fact that the NBBO shifts by \$0.05 on average results in an expected "mechanical" price impact that is higher in the nickel quoting regime than in the penny quoting regime (i.e., 0.08 x \$0.05 > 0.13 x \$0.016).

relative effective spread. As such, including it in the table provides a straightforward means to compare the average spread incurred by a liquidity-taking order (the relative effective signed spread) with the estimated profits available to the liquidity provider for that order (the realized spread).

[Insert Table 4]

In column 2, we present the average price impact for an observed trade after 1 microsecond for each of the three treatment groups as well as for the control group. All estimates are presented as the percentage (in basis points) of the NBBO midpoint at the microsecond of a trade. We first note that the control group mean as well as the estimated treatment effect (b_1-b_0) for each treatment group is non-zero, consistent with the mechanical effect of a trade dislodging the NBBO. Even if this effect is mechanical, it is worth noting that any such price impact remains a real barrier to a liquidity provider seeking to capture the wider spreads afforded by the tick size pilot. Reflecting this fact, average realized spreads in column 3 are generally lower than average relative effective spreads for both the control group and across treatment groups.

Somewhat surprisingly, point estimates for price impact at one microsecond are highest for TG3. We formally test the hypothesis of equal treatment effect, which we can reject at p<.01. This latter result suggests that, notwithstanding the much greater quoted depth for TG3 relative to TG1 and TG2, the greater depth of TG3 does not necessarily translate into lower expected price impact from a trade. In unreported tests, we evaluate the fraction of trades in TG3 in which a buy (sell) transaction results in a higher (lower) NBBO in the first microsecond following a trade, compared to the fraction of trades in TG1 and TG2. Overall, approximately 5.12% of trades within TG3 produce such a change in the NBBO, while the percentage for trades in TG1 an TG3 was 4.66% and 5.06%, respectively. As such, the slightly higher "mechanical" price impact for TG3 reflects the slightly large fraction of trades that dislodged the NBBO in TG3, notwithstanding the considerably greater depth created by the trade-at rule.

Columns 4 and 5 present average price impact and realized spreads as of the first second following an observed trade. By this point in time, average realized spreads remain significantly higher than those for the control group across all three treatment arms, consistent with the pilot's goal of enhancing liquidity provider profits through the wide tick size. However, realized spreads are also much lower than they were immediately following the trade. For example, average realized spreads for TG1 and TG2 are 33.6% and 35.7% lower at one second following a trade than at one microsecond after a trade. Realized spreads for TG3 have decayed at even greater rate, having declined by 42.6% over the same time period. Moreover, testing equality of treatment effect across each treatment group allows us to reject the possibility of equal treatment between TG3 and either TG1 or TG2 for both price impact and realized spreads. As with the higher "mechanical" effect on the NBBO for TG3, the greater decay rate for TG3 may be symptomatic of the trade-at rule: To the extent the rule forces more trading to occur on lit, public venues, a greater number of trades should result in greater price impact arising from both the mechanical effect noted previously as well as more trades being observed by fast market makers looking to adjust their quotes.[14]

Columns 6 and 7 provide average price impact and realized spreads, respectively, for the fifth minute following an observed trade. Notwithstanding the fact that the time span from an observed trade is 300 times as long as the time span utilized in columns 4 and 5, realized spreads have decayed by only an additional 38.0% for all three treatment groups—approximately the same rate that they decayed for TG1 and TG2 during the first second after a trade.

[14] In unreported results, we find that exchange trades in TG1 and TG have greater price impact than those reported in Table 4.

C. Trading Venues and Midpoint Trading

As noted in the Introduction, one of the more notable institutional features that can affect the relationship between wider ticks and liquidity provider profits is the fact that trading occurs across multiple, competing trading venues. Here, we evaluate how this competition among venues can interact with the pilot study's policy goal of subsidizing liquidity takers with wider tick sizes.

1. Liquidity Provider Profits and Exchange Pricing

We first evaluate the extent to which wider tick sizes shifted exchange trading from exchanges using maker/taker price schedules to those using taker/maker schedules. During our sample period, the majority of exchanges adhered to a "maker-taker" schedule, with a smaller group of exchanges adhering to an inverse "taker-maker" schedule that pays the rebate to liquidity takers, while charging a per-share fee to the liquidity provider. Two exchanges during our sample period (the Investors Exchange and the National Stock Exchange) either charged no fees at all or simply charged a small fee on all liquidity taking orders without paying a rebate.[15]

We present our results in Table 5. In Panel A, we first assess the extent to which the pilot study induced liquidity providers to transition away from maker-taker exchanges to taker-maker exchanges. We do so by calculating the fraction of the NBO for a security that was accounted for by the Best Offer (BO) posted on a particular exchange during the sample period. In column 1, we present the control group mean across all exchanges. Columns 2 through 4 present the mean for each treatment group. To facilitate comparison between maker-taker and taker-maker exchanges, we group exchanges separately according to whether they utilize a maker-taker schedule, a taker-maker schedule, or a zero rebate schedule. Overall, the first four columns

[15] The National Stock Exchange ceased trading on February 1, 2017 when it was acquired by the NYSE. Our trading data for this venue therefor runs from March 7, 2016 through January 31, 2017.

highlight a generate decline in the fraction of the NBO that is represented on maker-taker venues relative to taker-maker venues.

[Insert Table 5]

We formally examine the significance of this treatment effect in columns 5 through 7. Each column provides the regression coefficient for the estimated treatment effect (b_1-b_0) for each treatment group using the same regression framework used in Tables 4 and 5, including a control for the pre-period outcome. As shown in these columns, maker-taker venues showed a significance drop in the fraction of the NBO for a security that was accounted for by the venue's BO across all three treatment groups. The only exception was with respect to the NYSE MKT and the Chicago Stock Exchange, where their BOs made only a de minimis contribution to the NBO even in the case of control securities. Conversely, all three taker-maker venues experienced a significance increase in the fraction of the NBO that was represented by the BO on these exchanges. Indeed, across all three treatment arms, Nasdaq OMX saw its contribution to the NBO double from approximately 5% to 10%, while the mean for BATS Y rose from approximately 5% to approximately 13%. The two no-rebate exchanges likewise witnessed a general increase in the extent to which their BO contributed to the NBO across all three treatment arms. These findings are consistent with the tick size pilot inducing liquidity providers to compete on price by turning to taker-maker exchanges.

In Panel B, we present the same analysis but focus on the size-weighted fraction of all exchange trades that occurred on each exchange. Because we exclude non-exchange trades from the analysis, the panel reveals how liquidity providers' shift from quoting on maker-taker venues toward quoting on taker-maker venues translated into an equivalent shift in trading among those trades occurring on exchanges. Across all three treatment arms, the increase (decrease) in the fraction of the NBO that was represented by an exchange's BO during the pilot period was also reflected in an similar increase (decrease) in the fraction of exchange trades completed on that venue.

In light of the greater liquidity at the NBBO at taker-maker markets, these results are consistent with the incentive of liquidity taking orders to look for liquidity on taker-maker exchanges before maker-taker exchanges. They also underscore how the total profitability of a trade for a liquidity provider depends on the joint effects of the tick size and pricing schedule used among competing exchanges. For instance, ignoring other transaction fees, a "lead market maker" on BATS X who buys and sells passively at the NBBO in a control security with a penny quoted spread will earn $0.0129 per share ($0.01 of spread, plus rebates of $0.0045 x 2) by posting liquidity on BATS X given its current fee schedule. Conversely, the same market maker who makes such a pair of trades at the NBBO in a Pilot Security with a nickel quoted spread will earn $0.0464 ($0.05, less fees of $0.0018 x 2) by posting liquidity on BATS Y given its fee schedule. In this fashion, the interaction of the tick size pilot with exchanges' pricing schedules can reduce the effect of the subsidy on liquidity providers.

2. *Liquidity Provider Profits and Non-Displayed Order Types*

We examine how the pilot study affected the use of non-displayed liquidity—both with respect to the use of non-displayed orders on exchanges and the use of non-exchange venues—in Tables 6 and 7. As noted previously, prior research (Bartlett & McCrary, 2012; Kwan et al., 2014) suggests that the greater depth associated with wider tick sizes should induce liquidity providers to "queue jump" exchanges by posting non-displayed orders in non-exchange venues, generally offering price improvement over the NBBO. When the tick size is binding, Bartlett & McCrary (forthcoming) further demonstrate that SEC enforcement policy encourages such priceimproving orders to be "pegged" to the midpoint of the NBBO. However, this policy extends to exchange venues as well, and most exchanges permit traders to submit non-displayed orders that are pegged to the NBBO midpoint. Because exchanges observe price-time priority, pegged midpoint orders can permit traders to queue-jump the long queue lines whether submitted to an exchange or non-exchange venue. In short, the greater depth associated with the pilot study (especially in TG3) should encourage liquidity providers to "queue jump" by means of trading in non-exchange venues, as well as by using non-displayed midpoint orders on both exchange and non-exchange venues. At the same time, a liquidity provider who trades at the midpoint will by definition receive less than the quoted spread, underscoring how the exception for midpoint trades within the pilot study can work to reduce the subsidy to liquidity providers of nickel tick sizes.

In Table 6, we first evaluate how the tick size pilot affected the competition for order flow between exchange and non-exchange venues. In keeping with the format of Table 5, the first four columns present for the control group and each treatment group the mean size-weighted fraction of trades that occur on all exchanges and on all non-exchange venues. As shown in column 1, approximately 42% of all trades within the control group were made in non-exchange venues. This figure increases to approximately 46% and 45% respectively, for TG1 and TG2, respectively. In contrast, the figure declines to 32% for TG3, consistent with the trade-at rule reducing the number of trades that occur in non-exchange venues. We formally examine the significance of the treatment effect in columns 5 through 7, where we present the regression coefficient for the estimated treatment effect (b_1-b_0) for each treatment group using the same regression framework utilized in Tables 3-5. As shown in these columns, the foregoing changes in the fraction of trades occurring in non-exchange venues were statistically significant across all three treatment arms. In column 8, we additionally test whether the decline in non-exchange trading for TG3 was significantly different from the joint treatment effect in TG1 and TG2. As shown in the table, relative to the mean fraction of non-exchange trades in TG1 and TG2, the mean fraction of non-exchange trades in TG3 was approximately 13.4 percentage points lower, with the standard error indicating that the difference was statistically significant at $p<.01$.

[Insert Table 6]

Table 7 examines the incidence of midpoint trades within the sample, again using the same regression framework utilized in Tables 3-5. In columns 1 and 2, we first present the overall frequency of midpoint trades across both exchange and non-exchange venues. Column 1 indicates that the incidence of midpoint trades regardless of trading venue increased by approximately 5.7% for TG1, relative to a control group mean of 11.8%. The average treatment effect for TG2 was slightly higher at 5.9%; however, a formal test of equal treatment between TG1 and TG2 indicates that the difference is not statistically significant. However, as predicted, the average treatment effect was both higher for TG3 and statistically distinct from that of TG1 and TG2. As shown, the incidence of midpoint trades for TG3 nearly doubled relative to the control group. These results are largely unchanged in column 2 where we control for the pre-period frequency of midpoint trading.

[Insert Table 7]

In columns 3 and 4 we conduct the same analysis for all exchange trades, and in columns 5 and 6 we do so for all non-exchange trades. As shown in columns 3 and 4, the incidence of midpoint trades on exchanges increased across all three treatment arms; however, the magnitude of the effect was largest for TG1 and smallest for TG3. A likely explanation is that liquidity taking traders in TG3 securities were looking first to non-exchange venues for price improvement over the NBBO, reducing the probability that an in-bound marketable order would be placed on an exchange with a resting non-displayed midpoint order. Consistent with this claim, columns 5 and 6 highlight a notable increase in the frequency midpoint trades among non-exchange venues. Moreover, the sizes of the estimated treatment effects are consistent with each treatment group providing an incrementally greater incentive to post a pegged midpoint order in a non-exchange venue. The effect is especially notable for TG3 where the frequency of midpoint trades nearly triples, with an estimated two-thirds of all non-exchange trades taking the form of a midpoint execution. Given that nearly one-third of all trades in TG3 took place in a non-exchange venue, this finding helps explain why nearly one-third of trades in TG3 securities were filled in non-exchange venues despite the trade-at rule. The sizeable fraction of midpoint trades within TG3 also helps explain why, despite the trade-at rule, relative effective spreads were statistically indistinguishable across the three treatment groups.

VII. Conclusion

The tick size pilot represents a unique opportunity to study the efficiency with which a wider tick size creates a subsidy for liquidity provision in contemporary equity markets. While policy discussions over the past decade have commonly asserted that decimalization of stock prices in 2001 has harmed the incentive to provide liquidity in small capitalization equities, regulatory and technological changes since that time make it an open question whether reversing decimalization can produce liquidity benefits for smaller companies that outweigh the immediate increase in trading costs associated with wider ticks. As we show, the emergence of HFT market-making and the growing role of non-exchange trading can both interfere with the extent to which a wider tick size provides a subsidy for liquidity providers.

As noted by Rindi and Werner (2017), market maker profits have increased by approximately 40% in transactions involving securities that have been treated with nickel tick sizes. While this represents a notable increase in market maker profitability, our results underscore important institutional challenges facing the central policy goal of using wider tick sizes to subsidize liquidity providers. First, the exception for midpoint trades substantially reduces the aggregate subsidy available from wider tick sizes. Equally important, while liquidity taking trades in our sample nevertheless pay effective spreads that are on average more than twice what they were with a penny MPV, the sub-second decay rate of realized spreads makes the ability of liquidity providers to capture these spreads dependent on their capacity for trading at sub-second frequencies. Moreover, this challenge is made all the more acute to the extent trades are encouraged to occur on exchanges, making the relationship between the trade-at rule and the profitability of exchange market makers considerably more complicated than proponents of the trade-at rule have acknowledged. For instance, based on the estimates from Table 4, a liquidity provider who takes just one second to off-set a transaction will capture only one-third of the wider spreads for securities in TG1 or TG2. For trades in Treatment Group 3, such a liquidity provider will capture just over one-half of the wider trading costs.

Overall, these findings help explain why the profitability of making a market in treated securities is not higher than it is despite a quintupling of the tick size. For the same reason, they also point to the tick size as being a highly inefficient means to subsidize liquidity providers in contemporary equity markets.

References

Aït-Sahalia, Yacine, and Mehmet Saglam. 2017. High Frequency Market Making: Optimal Quoting. Working Paper, Princeton University.

Angel, James J., Lawrence E. Harris, and Chester S. Spatt. 2015. Equity Trading in the 21st Century: An Update. *Quarterly Journal of Finance* 5, 1-39.

Angel, James. J. 1997. Tick Size, Share Prices, and Stock Splits. *Journal of Finance* 52, 655-681.

Bacidore, Jeffrey, Robert Battalio, and Robert Jennings. 2003. Order Submission Strategies, Liquidity Supply, and Trading in Pennies on the New York Stock Exchange. *Journal of Financial Markets* 6, 337–362.

Bartlett, Robert P., III, and Justin McCrary. 2013 Shall We Haggle in Pennies at the Speed of Light or in Nickels in the Dark? How Minimum Price Variation Regulates High Frequency Trading and Dark Liquidity. Working paper, University of California, Berkeley School of Law.

Bartlett, Robert P., III, and Justin McCrary. 2017. How Rigged Are Stock Markets? Evidence from Microsecond Timestamps. Working paper, University of California, Berkeley School of Law.

Bartlett, Robert P., III, and Justin McCrary. Forthcoming. Dark Trading at the Midpoint: Pricing Rules, SEC Enforcement Policy, and Latency Arbitrage. *Journal of Law, Finance, and Accounting*.

Bessembinder, Hendrik. 2003. Trade Execution Costs and Market Liquidity After Decimalization. *Journal of Financial and Quantitative Analysis* 38, 747–777.

Bessembinder, Hendrik, and Kumar Venkataraman, 2010, Bid-Ask Spreads: Measuring Trade Execution Costs in Financial Markets, in "Encyclopedia of Quantitative Finance" (Rama Cont, Ed.), John Wiley & Sons.

Chakravarty, Sugato, Stephen Harris, and Robert Wood. 2001. Decimal Trading and Market Impact. Working paper, University of Memphis.

Comerton-Forde, Carole, Vincent Gregoire, and Zhuo Zhong. 2017. Inverted Fee Venues and Market Quality. Working paper, University of Melbourne.

Conrad, Jennifer, and Sunil Wahal. 2017. The Term Structure of Liquidity Provision. Working paper, University of North Carolina.

Coughenour J, Harris L. 2004. Specialist Profits and the Minimum Price Increment. Working paper, University of Delaware.

Cox, Justin, Bonnie Van Ness, and Robert Van Ness. 2017. Increasing the Tick: Examining the Impact of the Tick Size Change on Market Fee Venues. Working paper, University of Mississippi.

Gibson, Scott, Rajdeep Singh, and Vijay Yerramilli. 2003. The Effect of Decimalization on the Components of the Bid–Ask Spread. *Journal of Financial Intermediation* 12, 121–148.

Glosten, Lawrence R. and Paul R. Milgrom. 1985. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders, Journal of Financial Economics 14, 71-100.

Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka. 2009. Do Liquidity Measures Measure Liquidity? *Journal of Financial Economics* 92, 153-181.

Griffith, Todd, and Brian Roseman. 2017. Making Cents of Tick Sizes. Working paper, Utah State University Working Paper.

Harris, Lawrence. 1999. Trading in Pennies: A Survey of the Issues. Working paper, University of Southern California.

Holden, Craig, and Stacey Jacobsen. 2014. Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions, *Journal of Finance* 69, 1747-1785.

Huang, Roger D. and Hans R. Stoll. 1997. The Components of the Bid-Ask Spread: A General Approach. *Review of Financial Studies* 10, 995-1034.

Imbens, Guido. 2004. Nonparametric Estimation of Average Treatment Effects Under Exogeneity: A Survey, *Review of Economics and Statistics* 86, 4-30.

Hahn, Jinyong. 1998. On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects, *Econometrica* 66, 315-331.

Kwan, Amy, Ronald Masulis, and Thomas McInish. 2015. Trading rules, competition for order flows and market fragmentation. *Journal of Financial Economics* 115, 330-348.

Lee, Charles M.C. and Mark J. Ready, 1991, Inferring Trade Direction from Intraday Data, *Journal of Finance*, 46, 733-746.

Lin, Yiping, Peter Swan and Vito Mollica. 2017. Tick Size is Little More Than an Impediment to Liquidity Trading: Theory and Market Experimental Evidence. Working paper, University of New South Wales Working Paper.

Nallengara, Lona and John Ramsay. 2013. Testimony on "JOBS ACT Implementation Update", Before the Subcommittee on Investigations, Oversight and Regulations, Committee on Small Business, U.S. House of Representatives, available at https://www.sec.gov/news/testimony/2013-ts041113lnjrhtm.

Rindi, Barbara and Ingrid Werner. 2017. U.S. Tick Size Pilot. Working paper, Ohio State University Fish College of Business.

Ronen, Ravy, and Daniel Weaver. 2001. 'Teenies' Anyone? *Journal of Financial Markets* 4, 231–260.

Securities and Exchange Commission. 2015. Joint Industry Plans; Order Approving the National Market System Plan to Implement a Tick Size Pilot Program. Release no. 34-74892; File No. 4-657.

Wooldridge, Jeffrey M. 2010. Economic Analysis of Cross Section and Panel Data (2d. ed.)

Yao, Chen, and Mao Ye. 2014. Tick Size Constraints, Market Structure, and Liquidity. Working paper, University of Illinois at Urbana-Champaign.

Figure 1A. Incidence of Nickel−Priced Quotes by Treatment Assignment Over Time

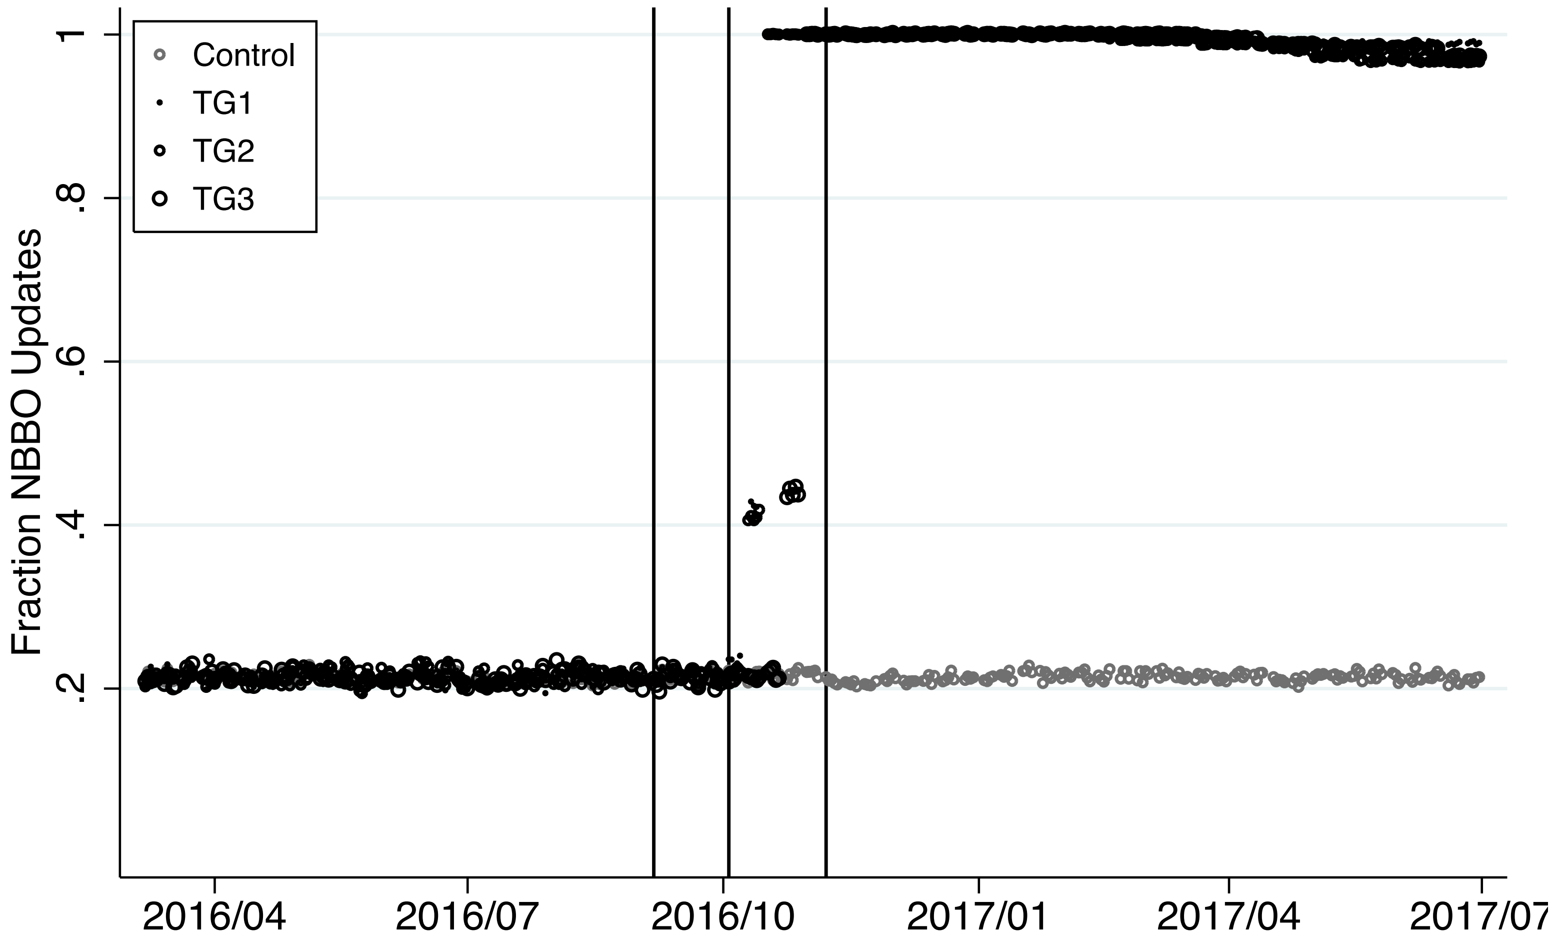

Figure 1B. Incidence of Nickel−Priced Trades
by Treatment Group Over Time

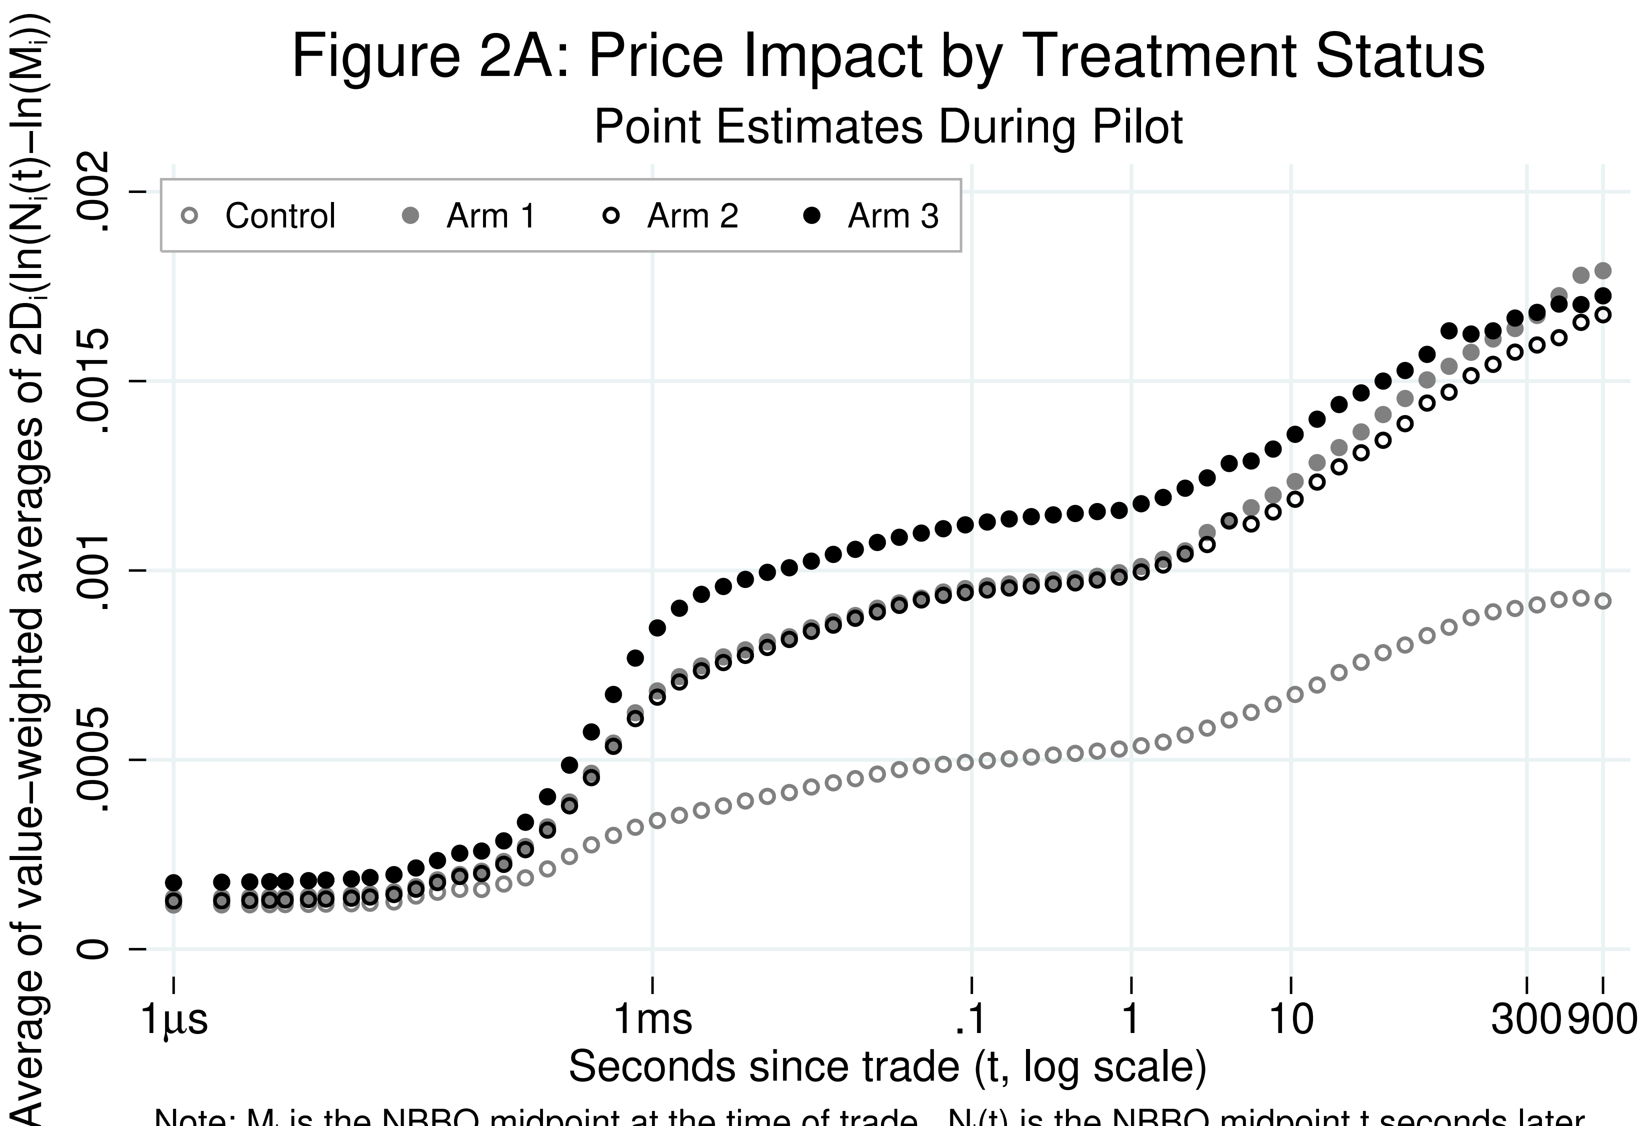

Note: M_i is the NBBO midpoint at the time of trade. $N_i(t)$ is the NBBO midpoint t seconds later.
See text for details.

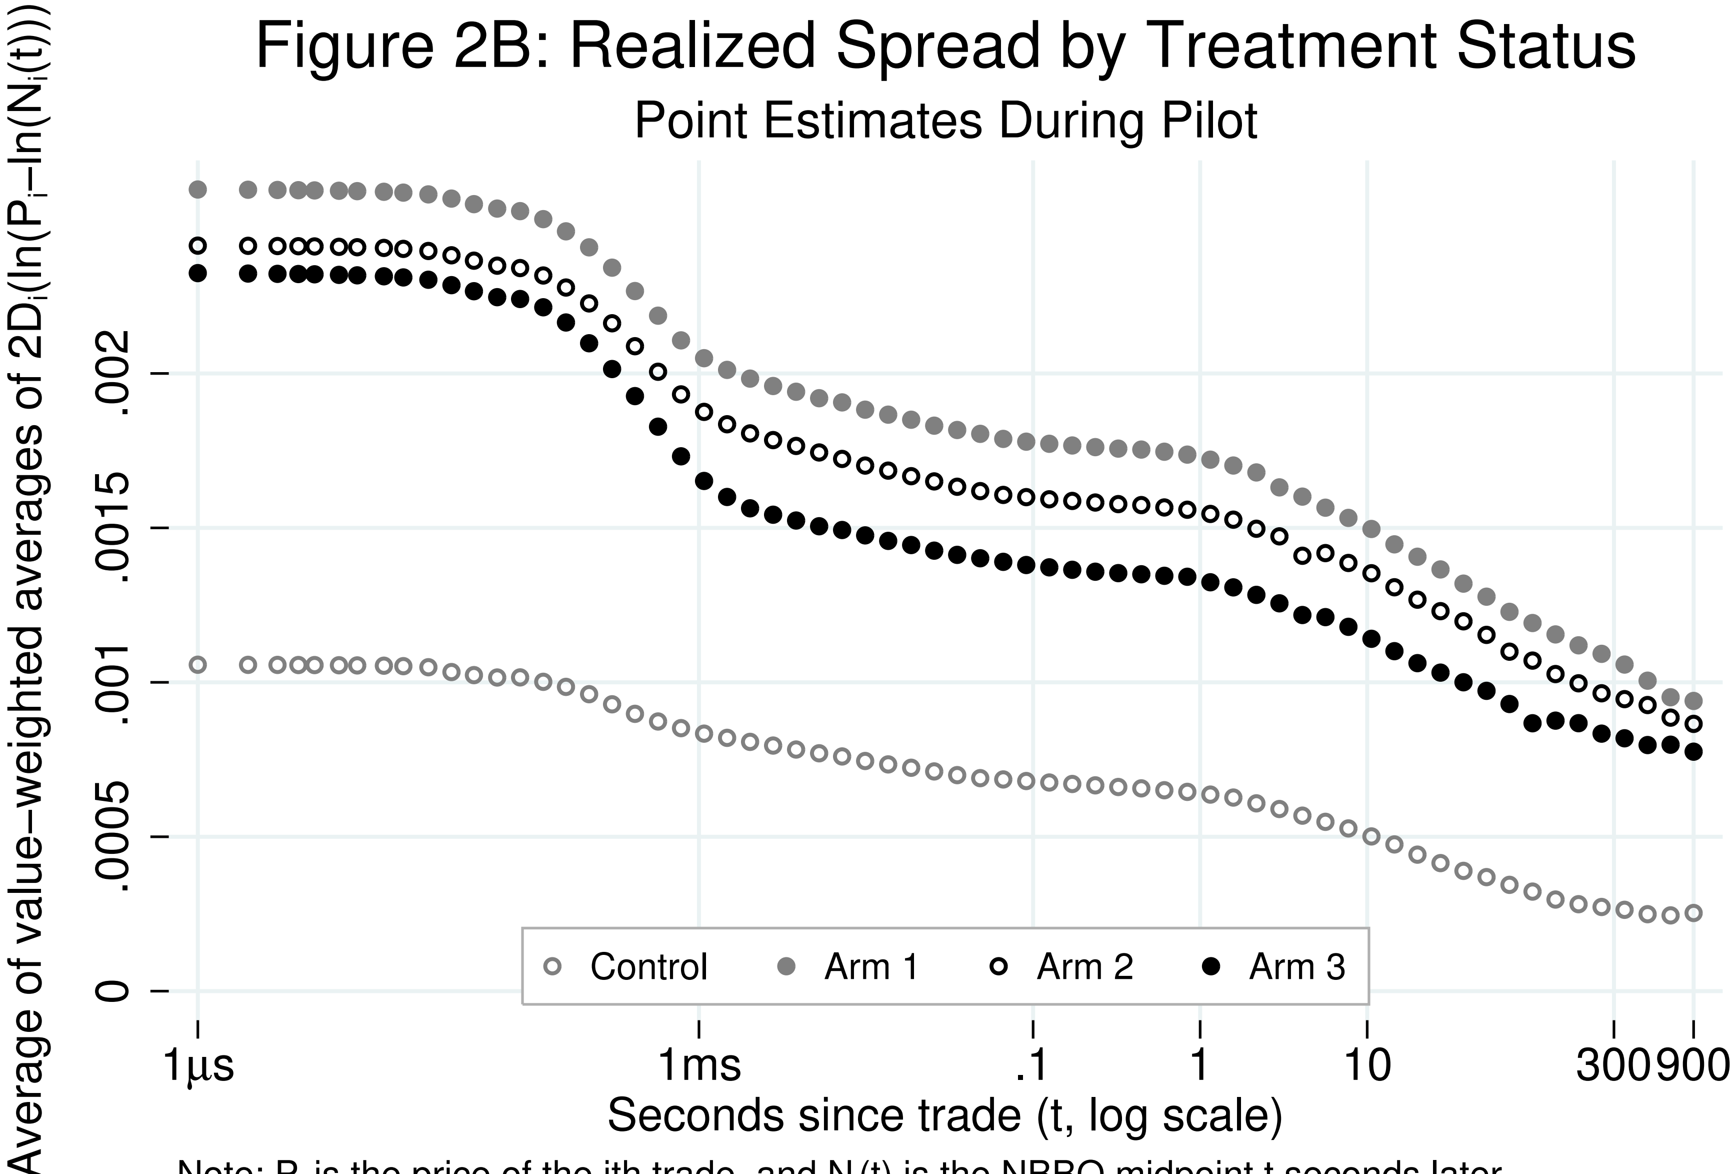

Note: P_i is the price of the ith trade, and $N_i(t)$ is the NBBO midpoint t seconds later.
See text for details.

Table 1: Average Quoted Spreads for Binding and Non-Binding Sample

This table examines the pilot study's average treatment effect on average quoted spreads (measured at the time of an observed trade) for Pilot Securities that had pre-pilot quoted spreads of less than $0.05 (Panel A) compared to Pilot Securities having pre-pilot quoted spreads of $0.05 or higher (Panel B). Each row represents a cross-sectional regression where each observation represents a size-weighted percent quoted spread for a given stock. That is, the outcome is the weighted average of the quoted-spread relative to the quote midpoint with weights of trade size, calculated separately for the four different time periods for a given stock. The outcome is multiplied by 100 for readability of coefficients. The four different time periods are described in detail in the text. Heteroskedasticity-consistent standard errors are presented in parentheses below point estimates. Number of observations is 1068 in panel A and 1100 in panel B, except for the Announcement Period where there are two stocks with no trades during the period and the number of observations is 1098. Regression specification is discussed in Section 5.

A. Stocks with Pre-Pilot Size-Weighted Quoted Spreads BELOW $0.05 (PRIMARY SAMPLE)

		Treatment Effect for Group:		
	Control Group	TG1	TG2	TG3
Pre-Pilot Period	0.3138	0.0037	-0.0296	0.0049
3/7/2016-9/2/2016	(0.0131)	(0.0263)	(0.0248)	(0.0287)
Announcement Period	0.3005	0.0132	-0.0354	-0.0129
9/6/2016-9/30/2016	(0.0144)	(0.0320)	(0.0265)	(0.0275)
Phase-In Period	0.3368	0.2413	0.1928	0.0973
10/3/2016-11/4/2016	(0.0175)	(0.0419)	(0.0411)	(0.0403)
Full Implementation	0.3145	0.3993	0.3166	0.3938
11/7/2016-6/30/2017	(0.0151)	(0.0541)	(0.0516)	(0.0559)

B. Stocks with Pre-Pilot Size-Weighted Quoted Spreads ABOVE $0.05 (EXCLUDED SAMPLE)

		Treatment Effect for Group:		
	Control Group	TG1	TG2	TG3
Pre-Pilot Period	1.2458	-0.0595	0.0069	0.0068
3/7/2016-9/2/2016	(0.0587)	(0.1195)	(0.1061)	(0.1108)
Announcement Period	1.5469	-0.5020	-0.2872	-0.4846
9/6/2016-9/30/2016	(0.3648)	(0.3746)	(0.3792)	(0.3741)
Phase-In Period	1.4024	-0.2298	-0.0994	-0.2799
10/3/2016-11/4/2016	(0.1050)	(0.1408)	(0.1471)	(0.1344)
Full Implementation	1.1111	-0.0546	0.0299	0.1132
11/7/2016-6/30/2017	(0.0577)	(0.1044)	(0.1047)	(0.1292)

Table 2: Tests of Validity of Randomization

This table reports balance tests of the 1068 control and treatment securities. Estimates the for Control Group Mean reflect regression-adjusted means after controlling for strata. Estimates for each Treatment Group reflect differences between the Control Group Mean and the Treatment Group Mean, controlling for strata. Standard errors are in parentheses.

		Treatment Group			
Outcome	Control Group Mean	1	2	3	Joint Test of Randomization
Log Trades	7.0297	0.0198	0.0240	-0.0469	0.8097
	(0.0207)	(0.0417)	(0.0418)	(0.0417)	p=0.4886
Log On-Exchange Trades	6.7286	0.0113	0.0235	-0.0595	0.9273
	(0.0223)	(0.0438)	(0.0439)	(0.0448)	p=0.4269
Log Off-Exchange Trades	5.6396	0.0358	0.0281	-0.0145	0.5479
	(0.0188)	(0.0395)	(0.0395)	(0.0374)	p=0.6496
Log Volume	11.9974	0.0144	0.0317	-0.0051	0.2394
	(0.0191)	(0.0436)	(0.0423)	(0.0386)	p=0.8689
Log On-Exchange Volume	11.4417	0.0036	0.0168	-0.0298	0.2840
	(0.0202)	(0.0446)	(0.0460)	(0.0408)	p=0.8370
Log Off-Exchange Volume	11.1006	0.0260	0.0506	0.0267	0.5528
	(0.0203)	(0.0461)	(0.0420)	(0.0403)	p=0.6463
Size-weighted ISO fraction	0.3198	-0.0051	-0.0072	-0.0084	2.3702
	(0.0020)	(0.0039)	(0.0037)	(0.0037)	p=0.0691
Quoted Spread	0.0268	0.0001	0.0000	-0.0007	0.3816
	(0.0005)	(0.0008)	(0.0009)	(0.0008)	p=0.7663
Percent Quoted Spread	0.2920	0.0027	0.0093	0.0162	0.5183
x100	(0.0072)	(0.0117)	(0.0144)	(0.0139)	p=0.6698
Locked or Crossed	0.0141	0.0004	0.0010	0.0002	0.4964
	(0.0004)	(0.0008)	(0.0008)	(0.0007)	p=0.6848
Effective Spread	0.0212	0.0000	-0.0001	-0.0006	0.3260
	(0.0004)	(0.0006)	(0.0007)	(0.0006)	p=0.8066
Relative Effective Spread	0.0023	0.0000	0.0001	0.0001	0.5658
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	p=0.6376
EQ Ratio	0.8436	-0.0034	0.0011	0.0027	1.0429
	(0.0014)	(0.0029)	(0.0029)	(0.0028)	p=0.3726
Price Improvement	0.0032	0.0001	0.0000	-0.0001	0.5932
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	p=0.6195
Log Inside Depth	1.5908	0.0105	0.0256	0.0499	0.6259
	(0.0191)	(0.0389)	(0.0387)	(0.0378)	p=0.5983

Table 3: Spreads and Depth

Panel A presents estimated treatment effects of the three pilot treatment arms on quoted spreads, relative quoted spreads, and relative effective spreads for sample securities. Panel B presents estimated treatment effects of the pilot on quoted inside depth, defined as one-half the sum of quoted depth at the national best bid and quoted depth at the national best ask. Panels present the estimated treatment as an additive effect to the Control Group Mean (b_1-b_0) and as a percentage of the Control Group Mean (b_1/b_0). Standard errors are in parentheses.

A. Spreads

	(1)	(2)	(3)	(4)	(5)	(6)
Estimated Effect	Quoted Spread		Relative Quoted Spread		Relative Effective Spread	
TG1						
b_1-b_0	0.0291	0.0288	0.4041	0.3979	0.3181	0.3169
	(0.0009)	(0.0008)	(0.0303)	(0.0281)	(0.0265)	(0.0258)
b_1/b_0	0.9492	0.9391	1.3305	1.2834	1.2889	1.2818
	(0.0478)	(0.0400)	(0.1153)	(0.1007)	(0.1234)	(0.1196)
TG2						
b_1-b_0	0.0292	0.0289	0.3764	0.3620	0.2822	0.2850
	(0.0011)	(0.0009)	(0.0347)	(0.0333)	(0.0289)	(0.0273)
b_1/b_0	0.9541	0.9422	1.2391	1.1674	1.1436	1.1530
	(0.0512)	(0.0432)	(0.1269)	(0.1158)	(0.1297)	(0.1224)
TG3						
b_1-b_0	0.0270	0.0274	0.4285	0.4124	0.2800	0.2777
	(0.0009)	(0.0009)	(0.0347)	(0.0312)	(0.0260)	(0.0246)
b_1/b_0	0.8831	0.8940	1.4107	1.3300	1.1346	1.1232
	(0.0466)	(0.0416)	(0.1291)	(0.1109)	(0.1192)	(0.1128)
Pre-Period Outcome		0.7019		0.9005		0.4314
		(0.0427)		(0.0956)		(0.1348)
Control Group Mean	0.0306	0.0307	0.3037	0.3101	0.2468	0.2472
	(0.0007)	(0.0006)	(0.0084)	(0.0068)	(0.0073)	(0.0070)
Observations	1068	1068	1068	1068	1068	1068
R^2	0.66	0.75	0.69	0.75	0.67	0.70
Control for Pre-period outcome	N	Y	N	Y	N	Y

B. Depth

Estimated Effect	(1)	(2)
	Inside Depth	
TG1		
b_1-b_0	1.4227	1.4160
	(0.0650)	(0.0522)
b_1/b_0	3.1483	3.1204
	(0.2698)	(0.2150)
TG2		
b_1-b_0	1.3592	1.3437
	(0.0630)	(0.0500)
b_1/b_0	2.8932	2.8331
	(0.2455)	(0.1916)
TG3		
b_1-b_0	1.6855	1.6285
	(0.0656)	(0.0534)
b_1/b_0	4.3950	4.0962
	(0.3539)	(0.2721)
Pre-Period Outcome		1.0092
		(0.0497)
Control Group Mean	1.7934	1.8073
	(0.0274)	(0.0201)
Observations	1068	1068
R^2	0.67	0.80
Control for Pre-period outcome	N	Y

Table 4: Price Impact and Realized Spreads

Table presents estimated treatment effects of the three pilot treatment arms on price impact and realized spreads at 1 microsecond after a trade, 1 second after a trade, and 5 minutes after a trade. Panels present the estimated treatment as an additive effect to the Control Group Mean (b_1-b_0) and as a percentage of the Control Group Mean (b_1/b_0). Standard errors are in parentheses.

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
			1 microsecond		1 second		5 minute	
Estimated Effect		Relative Effective Signed Spread (value weighted)	Price Impact	Realized Spread	Price Impact	Realized Spread	Price Impact	Realized Spread
TG1								
	b_1-b_0	0.1528	0.0025	0.1529	0.0486	0.1077	0.0770	0.0800
		(0.0001)	(0.0005)	(0.0132)	(0.0028)	(0.0119)	(0.0069)	(0.0089)
	b_1/b_0	1.3466	0.2171	1.4397	0.9131	1.6730	0.8441	2.9876
		(0.1253)	(0.0438)	(0.1443)	(0.0631)	(0.2173)	(0.0871)	(0.4993)
TG2								
	b_1-b_0	0.1367	0.0017	0.1361	0.0463	0.0914	0.0649	0.0697
		(0.0001)	(0.0005)	(0.0122)	(0.0026)	(0.0108)	(0.0066)	(0.0084)
	b_1/b_0	1.2047	0.1445	1.2818	0.8689	1.4196	0.7119	2.6027
		(0.1265)	(0.0438)	(0.1327)	(0.0591)	(0.1954)	(0.0818)	(0.4564)
TG3								
	b_1-b_0	0.1323	0.0056	0.1251	0.0621	0.0684	0.0745	0.0558
		(0.0001)	(0.0009)	(0.0107)	(0.0033)	(0.0090)	(0.0067)	(0.0074)
	b_1/b_0	1.1654	0.4886	1.1782	1.1654	1.0623	0.8165	2.0825
		(0.1176)	(0.0812)	(0.1190)	(0.0760)	(0.1620)	(0.0851)	(0.3940)
Pre-Period Outcome		0.4438	0.8556	0.4493	0.9068	0.3225	0.6715	0.1560
		(0.1391)	(0.1127)	(0.1223)	(0.0959)	(0.1005)	(0.0941)	(0.0620)
Control Group Mean		0.1135	0.0116	0.1062	0.0533	0.0644	0.0912	0.0268
		(0.0000)	(0.0003)	(0.0035)	(0.0012)	(0.0029)	(0.0026)	(0.0026)
Observations		1068	1068	1068	1068	1068	1068	1068
R^2		0.70	0.77	0.69	0.76	0.63	0.70	0.52
Control for Pre-period outcome		Y	Y	Y	Y	Y	Y	Y

Table 5: Exchange Trading

Panel A shows the fraction of the National Best Offer (NBO) that is accounted for by Best Offer (BO) posted on each exchange, as well as the estimated treatment effects of the three treatment arms on this measure. Panel B shows the size-weighted fraction of exchange trades for each exchange, as well as the estimated treatment effects of the three treatment arms on this measure. Estimated treatment effects represent coefficients on the treatment group indicators that are numerically equivalent to the efficient nonparametric estimator. Some figures may not add to 1 due to rounding. For all exchanges except for the Investors Exchange, estimates are based on control for the pre-period. Standard errors are in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Fraction of the NBO on a Given Exchange				Estimated Treatment Effects			
	Control	Treatment			TG1 vs.	TG2 vs.	TG3 vs.	TG3 vs.
Venue	Mean	TG1	TG2	TG3	Control	Control	Control	TG1 & TG2
A. Distribution of NBO Among Exchanges								
- Maker-Taker Venues:								
NYSE MKT (American) Stock Exchange	0.0085	0.0081	0.0077	0.0058	-0.0004	-0.0008	-0.0027	-0.0021
	(0.0004)	(0.0004)	(0.0006)	(0.0011)	(0.0006)	(0.0007)	(0.0012)	(0.0011)
New York Stock Exchange	0.1297	0.0957	0.0984	0.0916	-0.0340	-0.0313	-0.0382	-0.0055
	(0.0017)	(0.0027)	(0.0031)	(0.0030)	(0.0032)	(0.0036)	(0.0034)	(0.0036)
NYSE Arca SM	0.1357	0.1245	0.1265	0.1200	-0.0112	-0.0092	-0.0157	-0.0055
	(0.0014)	(0.0015)	(0.0018)	(0.0018)	(0.0020)	(0.0023)	(0.0023)	(0.0022)
NASDAQ Stock Exchange	0.3594	0.2694	0.2644	0.2595	-0.0899	-0.0950	-0.0999	-0.0074
	(0.0023)	(0.0031)	(0.0044)	(0.0034)	(0.0039)	(0.0050)	(0.0041)	(0.0043)
NASDAQ OMX PSX Stock Exchange	0.0198	0.0149	0.0145	0.0172	-0.0049	-0.0054	-0.0027	0.0025
	(0.0003)	(0.0004)	(0.0004)	(0.0005)	(0.0005)	(0.0004)	(0.0005)	(0.0005)
BATS Exchange	0.1021	0.0897	0.0880	0.0851	-0.0124	-0.0141	-0.0170	-0.0037
	(0.0010)	(0.0011)	(0.0015)	(0.0015)	(0.0015)	(0.0018)	(0.0018)	(0.0017)
Direct Edge X Stock Exchange	0.1149	0.1030	0.0993	0.0993	-0.0119	-0.0156	-0.0157	-0.0019
	(0.0011)	(0.0017)	(0.0022)	(0.0017)	(0.0020)	(0.0025)	(0.0020)	(0.0022)
Chicago Stock Exchange	0.0007	0.0007	0.0008	0.0008	0.0000	0.0000	0.0000	0.0000
	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
- Taker-Maker Venues:								
NASDAQ OMX BX Stock Exchange	0.0505	0.1091	0.1095	0.1283	0.0585	0.0590	0.0778	0.0190
	(0.0006)	(0.0014)	(0.0016)	(0.0020)	(0.0015)	(0.0017)	(0.0021)	(0.0022)
BATS Y-Exchange	0.0492	0.1283	0.1288	0.1306	0.0791	0.0796	0.0815	0.0021
	(0.0007)	(0.0014)	(0.0016)	(0.0014)	(0.0015)	(0.0017)	(0.0016)	(0.0018)
Direct Edge A Stock Exchange	0.0198	0.0349	0.0351	0.0351	0.0150	0.0152	0.0153	0.0001
	(0.0003)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0006)
- Zero Rebate Venues:								
Investors Exchange	0.0085	0.0269	0.0267	0.0256	0.0184	0.0182	0.0170	-0.0013
	(0.0002)	(0.0004)	(0.0004)	(0.0004)	(0.0005)	(0.0004)	(0.0005)	(0.0005)
National Stock Exchange	0.0001	0.0002	0.0002	0.0002	0.0001	0.0001	0.0001	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Size-Weighted Fraction of Trades				Estimated Treatment Effects			
	Control	Treatment			TG1 vs.	TG2 vs.	TG3 vs.	TG3 vs.
Venue	Mean	TG1	TG2	TG3	Control	Control	Control	Arms 1 & 2
B. Distribution of Trading Among Exchanges								
- Maker-Taker Venues:								
NYSE MKT (American) Stock Exchange	0.0066	0.0064	0.0061	0.0049	-0.0003	-0.0005	-0.0017	-0.0014
	(0.0004)	(0.0003)	(0.0004)	(0.0011)	(0.0005)	(0.0005)	(0.0012)	(0.0011)
New York Stock Exchange	0.1191	0.0859	0.0877	0.0820	-0.0333	-0.0314	-0.0371	-0.0048
	(0.0015)	(0.0026)	(0.0033)	(0.0030)	(0.0030)	(0.0036)	(0.0034)	(0.0037)
NYSE Arca SM	0.1204	0.0994	0.1005	0.0891	-0.0210	-0.0199	-0.0313	-0.0109
	(0.0015)	(0.0018)	(0.0016)	(0.0014)	(0.0023)	(0.0022)	(0.0021)	(0.0019)
NASDAQ Stock Exchange	0.3292	0.2556	0.2595	0.2400	-0.0736	-0.0698	-0.0892	-0.0175
	(0.0023)	(0.0031)	(0.0039)	(0.0032)	(0.0039)	(0.0045)	(0.0039)	(0.0041)
NASDAQ OMX PSX Stock Exchange	0.0106	0.0089	0.0084	0.0093	-0.0017	-0.0021	-0.0013	0.0007
	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
BATS Exchange	0.0915	0.0848	0.0828	0.0775	-0.0067	-0.0087	-0.0140	-0.0063
	(0.0008)	(0.0011)	(0.0013)	(0.0014)	(0.0014)	(0.0016)	(0.0016)	(0.0016)
Direct Edge X Stock Exchange	0.1370	0.1057	0.1043	0.0926	-0.0313	-0.0327	-0.0444	-0.0124
	(0.0013)	(0.0019)	(0.0020)	(0.0017)	(0.0023)	(0.0024)	(0.0022)	(0.0022)
Chicago Stock Exchange	0.0019	0.0022	0.0010	0.0014	0.0003	-0.0009	-0.0005	-0.0003
	(0.0004)	(0.0007)	(0.0002)	(0.0003)	(0.0009)	(0.0005)	(0.0005)	(0.0005)
- Taker-Maker Venues:								
NASDAQ OMX BX Stock Exchange	0.0543	0.1244	0.1245	0.1656	0.0701	0.0702	0.1113	0.0412
	(0.0005)	(0.0016)	(0.0019)	(0.0025)	(0.0017)	(0.0019)	(0.0026)	(0.0028)
BATS Y-Exchange	0.0603	0.1440	0.1421	0.1592	0.0837	0.0818	0.0989	0.0161
	(0.0006)	(0.0017)	(0.0018)	(0.0017)	(0.0018)	(0.0019)	(0.0018)	(0.0021)
Direct Edge A Stock Exchange	0.0268	0.0286	0.0286	0.0270	0.0018	0.0018	0.0002	-0.0016
	(0.0003)	(0.0005)	(0.0004)	(0.0004)	(0.0006)	(0.0005)	(0.0005)	(0.0005)
- Zero Rebate Venues:								
Investors Exchange	0.0415	0.0573	0.0559	0.0498	0.0157	0.0144	0.0083	-0.0068
	(0.0007)	(0.0015)	(0.0012)	(0.0011)	(0.0017)	(0.0013)	(0.0012)	(0.0014)
National Stock Exchange	0.0001	0.0001	0.0001	0.0001	0.0000	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)

Table 6: Exchange Vs. Non-Exchange Trading

Table shows size-weighted fraction of shares that were traded on all exchanges relative to all non-exchange trading venues during the sample period. Estimated treatment effects represent coefficients on the treatment group indicators that are numerically equivalent to the efficient nonparametric estimator. Some figures may not add to 1 due to rounding. For all exchanges except for the Investors Exchange, estimates are based on control for the pre-period. Standard errors are in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Size-Weighted Fraction of Trades				Estimated Treatment Effects			
	Control	Treatment			TG1 vs.	TG2 vs.	TG3 vs.	TG3 vs.
Venue	Mean	TG1	TG2	TG3	Control	Control	Control	TG1 & TG2
On Exchange	0.5822	0.5385	0.5506	0.6788	-0.0437	-0.0316	0.0966	0.1342
	(0.0022)	(0.0038)	(0.0045)	(0.0038)	(0.0044)	(0.0050)	(0.0044)	(0.0048)
Off Exchange	0.4178	0.4615	0.4494	0.3212	0.0437	0.0316	-0.0966	-0.1342
	(0.0022)	(0.0038)	(0.0045)	(0.0038)	(0.0044)	(0.0050)	(0.0044)	(0.0048)

Table 7: Midpoint Trading

Table presents estimated treatment effects of the three pilot treatment arms on the incidence of trades priced at the midpoint of the NBBO. Estimates are shown as an additive effect to the Control Group Mean (b_1-b_0) and as a percentage of the Control Group Mean (b_1/b_0). Standard errors are in parentheses .

	(1)	(2)	(3)	(4)	(5)	(6)
Estimated Effect	Overall		On Exchanges		Off Exchange Venues	
TG1						
b_1-b_0	0.057	0.056	0.053	0.053	0.034	0.033
	(0.003)	(0.002)	(0.003)	(0.003)	(0.004)	(0.003)
b_1/b_0-1	0.484	0.474	0.615	0.612	0.146	0.142
	(0.024)	(0.020)	(0.032)	(0.030)	(0.017)	(0.013)
TG2						
b_1-b_0	0.059	0.059	0.049	0.050	0.055	0.053
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.003)
b_1/b_0-1	0.502	0.499	0.567	0.574	0.232	0.225
	(0.022)	(0.021)	(0.027)	(0.027)	(0.018)	(0.015)
TG3						
b_1-b_0	0.095	0.095	0.039	0.039	0.444	0.447
	(0.003)	(0.003)	(0.002)	(0.002)	(0.008)	(0.006)
b_1/b_0-1	0.808	0.806	0.451	0.451	1.891	1.904
	(0.028)	(0.025)	(0.027)	(0.026)	(0.041)	(0.031)
Pre-Period Outcome		0.692		0.625		0.682
		(0.048)		(0.082)		(0.035)
Control Group Mean	0.118	0.118	0.086	0.086	0.235	0.235
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Observations	1068	1068	1068	1068	1068	1068
R^2	0.71	0.77	0.62	0.66	0.90	0.94
Control for Pre-period outcome	N	Y	N	Y	N	Y