



MIT Open Access Articles

Shades of darkness: A pecking order of trading venues

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Menkveld, Albert J., Bart Zhou Yueshen, and Haoxiang Zhu. "Shades of darkness: A pecking order of trading venues" <i>Journal of Financial Economics</i> vol. 24, no. 3 (2017): 503-534. © 2017.
As Published	http://dx.doi.org/10.1016/J.JFINECO.2017.03.004
Publisher	Elsevier BV
Version	Author's final manuscript
Citable link	https://hdl.handle.net/1721.1/124788
Terms of Use	Creative Commons Attribution-NonCommercial-NoDerivs License
Detailed Terms	http://creativecommons.org/licenses/by-nc-nd/4.0/

A version of this paper has been published in the Journal of Financial Economics.

"Shades of Darkness: A Pecking Order of Trading Venues." Menkveld, Albert J., Bart Zhou Yueshen, and Haoxiang Zhu. Journal of Financial Economics Vol.. 124, No. 3 (2017): 503-534.

DOI: 10.1016/j.jfineco.2017.03.004

Shades of Darkness: A Pecking Order of Trading Venues¹

Albert J. Menkveld²

Bart Zhou Yueshen³

Haoxiang Zhu⁴

First version: October 16, 2014

This version: June 26, 2016

¹For helpful comments and suggestions, we thank Bill Schwert (editor), an anonymous referee, Mark van Achter, Jim Angel, Matthijs Breugem, Adrian Buss, Sabrina Buti, Hui Chen, Jean-Edouard Colliard, Carole Comerton-Forde, John Core, Bernard Dumas, Thierry Foucault, Frank Hatheway, Terrence Hendershott, Pekka Hietala, Vincent van Kervel, Erik Loualiche, Ananth Madhavan, Katya Malinova, Jun Pan, Jonathan Parker, Christine Parlour, Joël Peress, Tālis Putniņš, Gideon Saar, Pedro Saffi, Astrid Schornick, Jeff Smith, Adrien Verdelhan, Jiang Wang, Avi Wohl, Hongjun Yan, Liyan Yang, Zhuo (Joe) Zhong, seminar/conference participants at VU University, MIT, and INSEAD, as well as Finance Down Under 2015, Second SEC Annual Conference on the Regulation of Financial Markets, The Financial Intermediation Research Society Conference 2015, 8th Annual SoFiE Conference, 7th Erasmus Liquidity Conference, European Financial Management Association 2015 Conference, the Society of Economic Dynamics Conference 2015, the China International Conference in Finance 2015, Inquire Europe 2015, and NBER Market Microstructure meeting 2015. We are grateful to NASDAQ for providing data. We also thank Sunny Li and Michael Abrahams for excellent research assistance. Menkveld gratefully acknowledges NWO for a VIDI grant.

²Vrije Universiteit Amsterdam and Tinbergen Institute; De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands; +31 20 598 6130; albertjmenkveld@gmail.com.

³INSEAD; Boulevard de Constance, 77300 Fontainebleau, France; +33 1 60 72 42 34; b@yueshen.me.

⁴MIT Sloan School of Management and NBER; 100 Main Street E62-623, Cambridge, MA 02142, USA; +1 617 253 2478; zhuh@mit.edu.

Shades of Darkness: A Pecking Order of Trading Venues

Abstract

We characterize the dynamic fragmentation of U.S. equity markets using a unique dataset that disaggregates dark transactions by venue types. The “pecking order” hypothesis of trading venues states that investors “sort” various venue types, putting low-cost-low-immediacy venues on top and high-cost-high-immediacy venues at the bottom. Hence, midpoint dark pools on top, non-midpoint dark pools in the middle, and lit markets at the bottom. As predicted, following VIX shocks, macroeconomic news, and firms’ earnings surprises, changes in venue market shares become progressively more positive (or less negative) down the pecking order. We further document heterogeneity across dark venue types and stock size groups.

Keywords: dark pool, pecking order, fragmentation

JEL Classifications: G12, G14, G18, D47

1 Introduction

A salient trend in global equity markets over the last decade is the rapid expansion of off-exchange, or “dark” trading venues. In the United States, dark venues now account for about 30% of equity trading volume (see Fig. 1(a) for an illustration of Dow-Jones stocks). In Europe, about 40% of volume trades off-exchange for leading equity indexes (see Fig. 1(b)).

Equally salient is the wide fragmentation of trading volume across dark venues. The United States has more than 30 “dark pools” and more than 200 broker-dealers that execute trades away from exchanges (see SEC, 2010). Dark pools, which are automated trading systems that do not display orders to the public, have grown fast in market shares and now account for about 15% of equity trading volume in the U.S., according to industry estimates.¹ In Europe, dark venues also face a high degree of fragmentation, with at least ten multilateral dark venues operating actively.²

The fragmentation of trading—between exchanges and dark venues, as well as across dark venues—is a double-edged sword. It creates a conflict between the efficient interaction among investors and investors’ demand for a diverse set of trading mechanisms. The SEC (2010, p. 11-12) highlights this tradeoff in its Concept Release on Equity Market Structure:

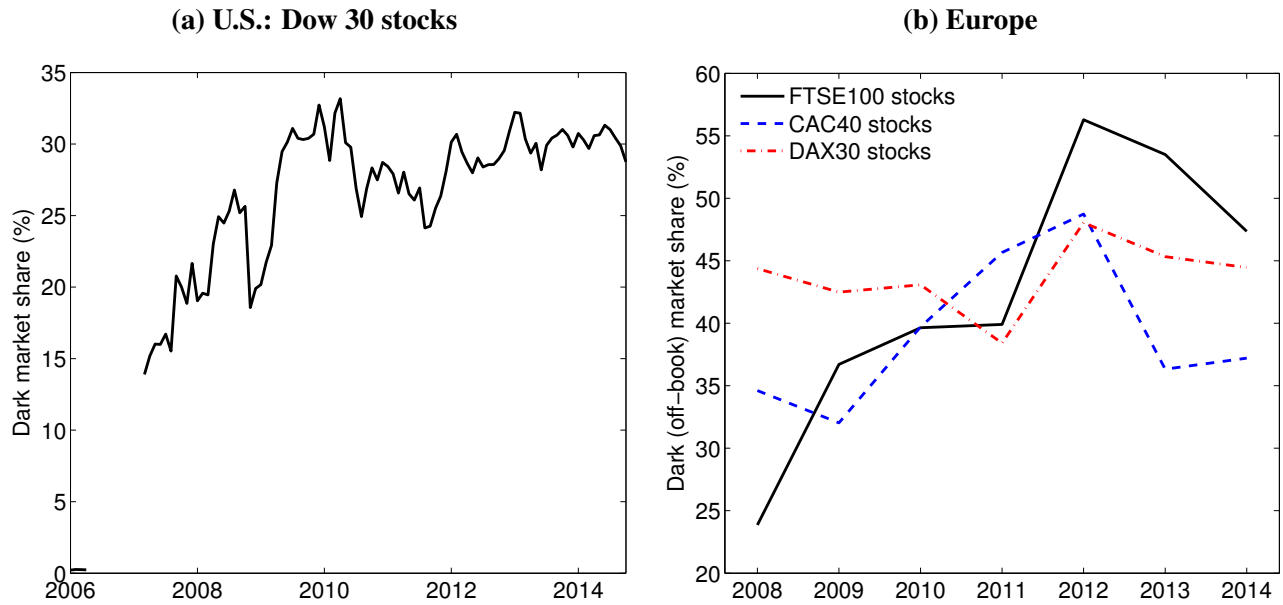
Fragmentation can inhibit the interaction of investor orders and thereby impair certain efficiencies and the best execution of investors’ orders. . . . On the other hand, mandating the consolidation of order flow in a single venue would create a monopoly and thereby lose the important benefits of competition among markets. The benefits of such competition include incentives for trading centers to create new products, provide high quality trading services that meet the needs of investors, and keep trading fees low.

An important step toward a better understanding of this tradeoff is to empirically characterize the degree of heterogeneity among trading venues. If venues appear homogeneous, the “liquidity-begets-liquidity” insight from early theories (see, for example, Pagano, 1989 and Chowdhry and Nanda, 1991) would imply that fragmentation is generally harmful; and that the fragmentation of dark venues, which provide little or no pre-trade transparency, causes a particular concern because investors cannot observe the presence of counterparties *ex ante* and must engage in costly search in multiple dark venues. In contrast, if venues do demonstrate heterogeneity, and if theory provides an economic rationale for such behavior, then fragmentation could be viewed as a natural equilibrium outcome and not necessarily a concern. For example, recent theories of dark pools show that precisely because of

¹Industry estimates are provided by Tabb Group, a consultancy firm, and Rosenblatt Securities, a broker. On June 2, 2014, FINRA started publishing weekly statistics of transaction volumes in alternative trading systems (ATS), with a two-week lag. Many dark pools are registered as ATS. For more details, see <https://www.finra.org/Newsroom/NewsReleases/2014/P519139>.

²See “European Dark Trading Analysis,” Fidessa, October 2013.

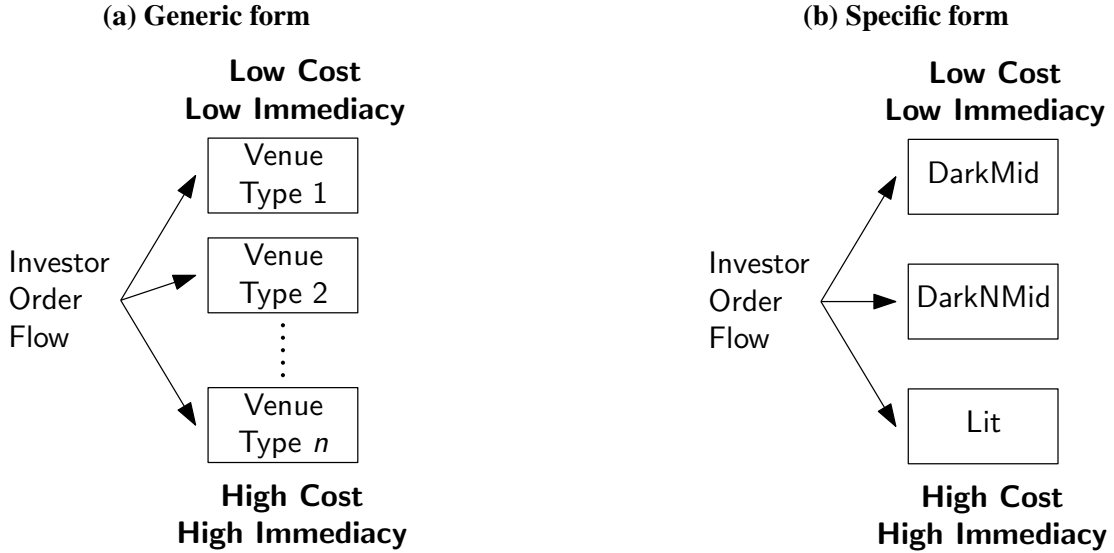
Figure 1: Dark market share in United States and Europe. This figure shows the market shares of dark trading in the U.S. and in Europe. Panel (a) plots the monthly average dark shares of the 30 stocks in the Dow Jones Industrial Average from 2006 to October 2014. We use the stocks that are in the Dow index as of November 2014. Volume data are obtained from Bloomberg and TAQ. Dark trades are defined by those reported to FINRA (code “D” in TAQ definition). From May 2006 to February 2007 the estimates are missing because TAQ data mix trades reported to FINRA with some NASDAQ trades. Panel (b) plots the averages of dark shares of FTSE100, CAC40, and DAX30 index stocks. These estimates are directly obtained from Fidessa.



their pre-trade opacity and the associated execution uncertainty, dark venues attract a different type of investors from those on the exchanges (Hendershott and Mendelson, 2000; Degryse et al., 2009; Buti et al., 2015, Ye, 2011; Zhu, 2014; Brolley, 2014). Under this “venue heterogeneity” or “separating equilibrium” view, fragmentation is an equilibrium response to the heterogeneity of investors and time-varying market conditions.

Pecking order hypothesis. In this paper we characterize the dynamic fragmentation of U.S. equity markets. Through the lens of dynamic fragmentation we gain insights into the degree of heterogeneity among trading venues and hence the important tradeoff highlighted in the SEC remark. In particular, we propose and test a “pecking order hypothesis” (POH): when executing orders investors “sort” dark and lit venues by the associated costs (bid-ask spread, price impact, information leakage) and immediacy, in the form of a “pecking order.” The top of the pecking order consists of venues with the lowest cost and lowest immediacy, and the bottom of the pecking order consists of venues with the highest cost and highest immediacy. The pecking order hypothesis predicts that, as investors’ trading needs become more urgent, they move from low-cost, low-immediacy venues to high-cost,

Figure 2: Pecking order hypothesis. This figure depicts the pecking order hypothesis. Panel (a) shows the generic form. Panel (b) shows the specific form. Midpoint crossing dark pools (DarkMid) are on top. Non-midpoint dark pools (DarkNMid) are in the middle. Lit markets (Lit) are at the bottom. Detailed descriptions of various dark pool types are collated in Section 3.1.



high-immediacy venues. This intuitive sorting is illustrated in Panel (a) of Fig. 2.

More concretely, recent theories of dark pools mentioned above predict that dark venues are at the top of the pecking order, whereas lit venues are at the bottom, because dark venues provide potential price improvement but lack execution certainty. In addition, we rank two important categories of dark pools—midpoint dark pools (“DarkMid”) and non-midpoint dark pools (“DarkNMid”)—according to the extent to which prices are constrained in these venues. Because midpoint dark pools restrict transaction prices to the midpoint of the national best bid and offer (NBBO), which confines market clearing, these dark pools offer the highest potential cost saving but the lowest immediacy. In contrast, non-midpoint dark pools are effectively non-displayed limit order books that allow transactions anywhere within the NBBO. As we show in a simple model, the cost of trading in these venues is between that on exchanges and that in midpoint dark pools; and so is the immediacy. This specific ordering of the three venue types is illustrated in Panel (b) of Fig. 2. This pecking order captures both “exchanges are liquidity of last resort” and “not all dark pools are created equal.”

We formally develop the specific-form pecking order hypothesis in a stylized model and thus provide some micro-foundation for it. The analysis shows that the pecking order obtains in a Bayesian-Nash equilibrium where large, liquidity-driven investors minimize transaction cost when hit by exogenous trading needs. It adds to existing models by focusing on the three-way fragmentation (DarkMid, DarkNMid, and Lit) as opposed to the standard dark-lit split in previous studies.

We test the pecking order hypothesis in U.S. equity market by exploring a unique dataset on dark trading. Our dataset disaggregates dark transactions into five categories by trading mechanism, including the two types of dark pools shown in Panel (b) of Fig. 2. The other three categories of dark transactions are retail investors' trades internalized by broker-dealers, average-price trades, and other (mostly institutional) trades. The detailed descriptions are provided in Section 3. To the best of our knowledge, this dataset provides the most comprehensive and granular view of U.S. dark trading that is accessible by academics. Without a disaggregated dataset like this, testing the pecking order hypothesis—or conducting any analysis of dark venue heterogeneity—would be very difficult because all off-exchange trades in the U.S. are reported under a single consolidated category called “trade reporting facilities” (TRFs). Data aggregation at TRFs masks the very heterogeneity in trading mechanisms that we are interested in.

We estimate a panel vector-autoregressive model with exogenous variables to characterize the dynamic interrelation among dark volumes, endogenous market conditions, and exogenous shocks (see Section 4). Key to our empirical strategy are three exogenous variables: VIX, macroeconomic data releases, and earnings surprise. We use these variables as proxies for shocks to investors' demand for immediacy. Applied to this setting, the pecking order hypothesis predicts that the proportional changes in venue market shares after these shocks should become progressively more positive (or less negative) the further down in the pecking order.

The data support the pecking order hypothesis. Following a 0.01% upward shock to VIX, dark pools that cross orders at the midpoint lose 4.6% of their market share (from 2.4% to 2.29%), dark pools that allow some price flexibility lose 3.3% of their market share (from 7.52% to 7.27%), and lit venues gain 1.0% of their market share (from 76.2% to 77.0%).

The effect is much larger for macroeconomic data releases. In the minute right after the macroeconomic news, DarkMid and DarkNMid market shares are 38.5% and 22.7% lower than their steady state levels, but Lit market share is 9.1% higher than its steady state level. The same qualitative pattern is observed after surprise earnings, although there only the DarkMid change is statistically significant. The shifts of market shares are not due to reduced trading volumes in dark venues after urgency shocks; instead, trading volumes in all venues increase but the increase in lit venues significantly outweighs that in dark venues.

The stylized model that delivers the pecking order hypothesis also predicts that the pecking order should be more evident in low-volume or high-spread stocks, (i.e., illiquid stocks). We verify this heterogeneity by running the same VARX model on large, medium, and small stocks separately. In these subsamples, the pecking order hypothesis finds strong support in medium and small stocks, but has no statistical significance for large stocks. This confirms the model's prediction.

We believe the interpretation of our results warrant a couple of final remarks. First, while our

results suggest that the fragmentation among dark venue types can be an equilibrium response to investor heterogeneity and time-varying market conditions, our analysis is silent on the fragmentation within each dark venue type. The latter question requires more detailed data on venue identities, not only venue types. Second, the pecking order hypothesis implicitly assumes that at least some investors make rational venue choices based on correct information of how these venues operate. This point is important in light of recent cases that certain dark pools are charged with misrepresenting or failing to disclose information to investors.³

Related literature. The primary contribution of this paper to the literature is the characterization of *dynamic* fragmentation of dark and lit venues through the pecking order hypothesis. Rather than assessing the market composition from a static view, we study how market shares evolve dynamically upon certain urgency shocks. Our approach of focusing on exogenous shocks (VIX, macroeconomic news, and surprise earnings) differs from most existing empirical studies on dark venues, which typically relate dark trading to endogenous measures of market quality (e.g., spread, depth, and volatility).

Access to unique, comprehensive, and granular data on various types of dark trading enabled us to study the pecking order hypothesis. Most existing studies of dark venues use either aggregate off-exchange trades or trades in a few selected dark pools. Studies based on aggregate dark trades include O’Hara and Ye (2011), Hatheway, Kwan, and Zheng (2013), and Degryse, de Jong, and van Kervel (2015a). Studies based on trades in a few dark pools include Hendershott and Jones (2005) (Island ECN), Ready (2014) (Liquidnet and POSIT), Buti, Rindi, and Werner (2011) (11 anonymous dark pools), Boni, Brown, and Leach (2012) (Liquidnet), Nimalendran and Ray (2014) (one anonymous dark pool), and Foley, Malinova, and Park (2013) (dark order types on Toronto Stock Exchange).

Equipped with more granular data, a few recent studies have devoted attention to dark venue heterogeneity. The vast majority of them use non-U.S. data. In Australian equity market, Comerton-Forde and Putniņš (2015) find that block dark trades and non-block dark trades on Australian Securities Exchange have different implications for price discovery, as measured by autocorrelation, variance ratio, or short-term return predictability. Three studies on the Canadian equities explore the asymmetric effects of a new “trade-at” regulation on midpoint and non-midpoint executions.⁴ Foley

³For example, in October 2011, SEC finds that Pipeline, a dark-pool operator who claimed to only allow buy-side firms to participate, had filled the majority of customer orders through its own trading affiliate (see <http://www.sec.gov/news/press/2011/2011-220.htm>). On June 25, 2014, Eric Schneiderman, Attorney General of the State of New York, alleges that Barclays falsified marketing material and misrepresented information to clients about the presence of high-frequency traders in its dark pool. In August 2015, SEC charges ITG, a dark pool operator, with operating secret trading desk and misusing dark pool subscriber trading information (see <http://www.sec.gov/news/pressrelease/2015-164.html>).

⁴A trade-at rule stipulates that when a marketable order arrives, a venue that does not already offer the best quote must provide sufficiently large price improvement, unless the order is sufficiently large. The minimum price improvement mandated by the Canadian trade-at rule is one tick, unless the bid-ask spread is one tick, in which case the minimum price

and Putniņš (2015) conclude that market quality deteriorates after the trade-at rule, but Comerton-Forde, Malinova, and Park (2015) and Devani, Anderson, and Zhang (2015) counter that little deterioration is detected in their own analyses. Comerton-Forde, Malinova, and Park (2015) in particular argue that changes in the Canadian dark market share after trade-at rule can be attributed to a single venue that takes retail orders. In Dutch equity markets, Degryse, Tombeur, and Wuyts (2015b) find that hidden orders on exchanges and off-exchange dark trades are affected differently by market conditions such as volume, spread, and depth. It should be emphasized that the market structures in Australia, Canada, and the Netherlands differ significantly from that of the U.S. and conclusions from these studies cannot literally be applied to U.S. equity markets.

The only other academic paper we are aware of that has a similar dataset for U.S. equities is Kwan, Masulis, and McNish (2015). They examine how the minimum tick size affects competitiveness of exchanges relative to dark venues, which is a research question very different from ours. Using a regression discontinuity design around the \$1.00 price threshold, they conclude that the tick size constraint weakens the competitiveness of exchanges.

2 A pecking order hypothesis of trading venues

In this section, we further motivate the pecking order hypothesis in its specific form as in Panel (b) of Fig. 2: Dark pools that cross orders at the midpoint of NBBO (labeled “DarkMid”) are on the top of the pecking order, dark pools that allow some price flexibility (labeled “DarkNMid”) are in the middle, and transparent venues (labeled “Lit”) are at the bottom. It is this specific-form hypothesis that will be the main hypothesis discussed and tested throughout the remainder of the paper.

The ordering that midpoint dark venues sit on top of lit venues is predicted by existing theories of dark pools. For example, Hendershott and Mendelson (2000), Degryse, Van Achter, and Wuyts (2009), and Zhu (2014) all model the competition between an exchange (that has dedicated liquidity providers) and a midpoint dark pool. They predict that because of execution uncertainty in dark pools, investors are more likely to send orders to the lit venue if urgency goes up—driven by either a higher delay cost for unfilled orders or a higher value of proprietary information. Buti, Rindi, and Werner (2015) model the competition between a limit order book and a midpoint dark pool. They show that dark pool market share is higher if the order book is more liquid (i.e., have longer queues). To the extent that a more liquid book reduces the opportunity cost of unfilled orders, their prediction is broadly consistent with the other papers.

DarkNMid is a “lighter shade of dark,” sitting between DarkMid and Lit. In practice, DarkNMid often operates as nondisplayed limit order books, in which transactions can happen anywhere between

improvement is a half tick.

the national best bid and national best offer. The execution price can respond to supply and demand; meanwhile, the trade-through restriction implies that investors still get a price improvement relative to Lit. In other words, the cost and benefit of using DarkMid, rather than Lit, should apply to DarkNMid as well, albeit to a less degree. For example, the online appendix of Zhu (2014) provides a model of a non-midpoint dark pool, which runs as a uniform-price divisible auction, but with rationing of orders whenever the dark pool price hits the constraint imposed by a “trade-at” rule. He shows that price flexibility in non-midpoint dark pools weakens its effectiveness in filtering out informed orders.

According to the pecking order hypothesis, upon urgency shocks, DarkMid and DarkNMid should lose market share to Lit, and the drop in DarkMid market share should be particularly large. Section 3.4 discusses our empirical proxies for urgency. Section 8 proposes a simple model to micro-found our pecking order hypothesis.

Before describing the data and conducting empirical analysis, a few comments and clarifications are in order. First, the pecking order hypothesis is a dynamic one. It says that after investors receive urgency shocks in real-time, volume shares should become progressively larger down the pecking order. This implies that many stable or slow-moving determinants of market shares, such as trading fees and membership restrictions, are orthogonal considerations.

Second, the pecking order hypothesis applies to both aggressive (market order) and passive (limit order) execution strategies of end investors. After all, each transaction needs a buyer and a seller. As investors who use aggressive strategies move down the pecking order, ready to pay a higher cost, investors using passive strategies would also move down the pecking order to meet such demand. If broker-dealers and high-frequency trading firms can also participate in dark venues, their location of liquidity provision should follow the flow of end investors as well.

Third, the specific form of pecking order hypothesis builds on the key features that dark-pool prices are constrained by lit quotes and that dark execution is uncertain. Opacity (“darkness”) is essential, for otherwise the very purpose of not publicly disclosing trading interest is defeated. Rationing and price constraints are also essential, for otherwise execution risk is much smaller. One could think of an alternative setting with two competing limit order books, one lit and the other dark, and the prices in both can freely move. The pecking order hypothesis cannot be applied to such a setting, as the trade-through restriction is not satisfied. In other words, if considered outside the market reality of the U.S. equity market, opacity alone is not sufficient to generate the pecking order hypothesis.

Fourth, the pecking order hypothesis is about the coexistence of venues. It does not rank market structures that have one dominant venue. For example, in an extension of Kyle (1989) with informed liquidity provision, Boulatov and George (2013) compare a lit-only market with a dark-only market, and finds that the opaque market offers better price discovery as it encourages more informed traders to be liquidity providers. In an experimental setting, Bloomfield, O’Hara, and Saar (2015) show that

although traders' strategies are greatly affected by the degree of opacity, market outcomes are largely invariant to opacity. Insights from these venues are complementary to ours.

Fifth, the cost-versus-immediacy tradeoff underlying the pecking order hypothesis is similar in spirit to the tradeoff between market and limit orders in a centralized limit order book setting. Parlour and Seppi (2008) survey the limit order book literature, both theory and empirics. Given today's fragmentation and heterogeneous trading mechanisms, we believe that *venue choice* is an important dimension to study over and above the *order type choice* studied in the limit order book literature.

Sixth, we interpret the urgency shocks underlying the pecking order hypothesis as primarily due to liquidity needs. Our empirical proxies for urgency discussed in Section 3.4 are, generally, public information releases, and our suggestive model of Section 8 assumes symmetric information about asset fundamentals. Although asymmetric information is not explicitly dealt with in this paper, the pecking order hypothesis is also consistent with existing dark pool models with asymmetric information. For instance, in a model that combines both asymmetric information and liquidity needs, Zhu (2014) finds that the market share of the dark pool tends to decrease in the value of the proprietary information.

Lastly, there is an alternative motivation for the pecking order hypothesis, based on an agency conflict between investors and their brokers. Its starting point is that brokers decide where to route investors' orders, and investors monitor brokers insufficiently. If brokers earn more profits by routing investors' orders first to their own dark pools, then orders would first flow into broker-operated dark pools and then to other dark and lit venues. Only when investors tell brokers to execute quickly will they have no choice but to send it to lit venues. This alternative motivation is less rich as it does not suggest a particular ranking of dark venue types. It simply states that dark venues take priority, whatever type of dark venue the broker is running, except when investors demand quick execution.

3 Data and summary statistics

Our data sample covers 117 stocks in October 2010 (21 business days). It consists of large-cap, medium-cap, and small-cap stocks in almost equal proportions.⁵ This section presents the various data sources, introduced sequentially in the first four subsections. Each subsection also describes and motivates the model variables that are based on them. The section concludes with presenting summary

⁵This sample of stocks is the same as that in the NASDAQ HFT dataset used by many studies of HFT. The original sample compiled by NASDAQ contains 120 stocks, with 40 stocks in each size category. But only 117 of the 120 stocks are present in our sample period. Brogaard, Hendershott, and Riordan (2014) state that "The top 40 stocks are composed of 40 of the largest market capitalization stocks. . . The medium-size category consists of stocks around the 1000th largest stocks in the Russell 3000. . . , and the small-size sample category contains stocks around the 2000th largest stock in the Russell 3000."

statistics based on the final data sample.

3.1 Dark volumes

In the United States, off-exchange transactions in all dark venues are reported to trade-reporting facilities (TRFs). The exact venue in which the dark trade takes place is not reported in public data. More recently, FINRA started to publish weekly transaction volumes in alternative trading systems, but these volumes are not on a trade by trade basis.⁶ The TRF managed by NASDAQ is the largest TRF, accounting for about 92% of all off-exchange volume in our sample. Our first dataset contains all dark transactions reported to the NASDAQ TRF. These trades are executed with limited pre-trade transparency.⁷ A salient and unique feature of our data is that the dark transactions are disaggregated into five categories by trading mechanism. The exact method of such disaggregation is proprietary to NASDAQ, but we know their generic features. The five categories include:

1. **DarkMid.** These trades are done in dark pools that use midpoint crossing as much as possible. Midpoint crossing means that the buyer and the seller in the dark venue transact at the midpoint of the National Best Bid and Offer (NBBO). “Agency-only” dark pools (i.e., those without proprietary order flow from the dark pool operators) typically work this way.
2. **DarkNMid.** These trades are done in dark pools that allow flexibility in execution prices (not necessarily midpoint). This feature leads us to believe that dark pools operated by major investment banks belong to this category.
3. **DarkRetail.** These trades come from retail investors and are internalized by broker-dealers. Retail brokers often route order flow submitted by retail investors to major broker-dealers, who then fill these orders as principal or agent. These transactions would be classified as DarkRetail.
4. **DarkPrintB.** These trades are “average-price” trades. A typical example is that an institutional investor agrees to buy 20,000 shares from a broker, at a volume-weighted average price plus a spread. This trade of 20,000 shares between the investor and the broker would be classified as a “print back” trade, abbreviated as “PrintB.”

⁶See <https://www.finra.org/Newsroom/NewsReleases/2014/P519139>.

⁷Our dark transaction data do not include trades on electronic communication networks (ECNs). ECNs are transparent venues that register as alternative trading systems (ATS), but they are not exchanges for regulatory purposes. In our sample, ECNs account for a very small fraction of total transaction volume. BATS and DirectEdge, two major exchanges that recently merged, used to be ECNs, but they have converted to full exchange status in November 2008 and July 2010, respectively.

Table 1: Volume shares of dark venues and the aggregate lit venue. This table shows the average volume shares of various venue types. Each venue share is computed as the percentage of the volume executed in that venue over the total volume obtained from TAQ in our sample period.

	full sample	large	medium	small
<i>DarkMid</i> [percent]	2.48	2.47	2.59	2.62
<i>DarkNMid</i> [percent]	9.08	9.20	8.11	6.55
<i>DarkRetail</i> [percent]	12.86	13.19	7.76	11.29
<i>DarkPrintB</i> [percent]	0.99	1.02	0.47	0.65
<i>DarkOther</i> [percent]	6.81	6.79	6.29	9.17
<i>Lit</i> [percent]	67.79	67.32	74.78	69.72

5. **DarkOther.** These are dark trades not covered by the categories above. A typical example in this category would be a negotiated trade between two institutions on the phone (i.e., not done on any electronic platform).

We emphasize that each category is not a single trading venue, but a collection of venues that are qualitatively similar in terms of their trading mechanism. In the interest of brevity, however, we will use the terms “venue” and “type of venue” interchangeably. For concreteness, Appendix A presents a snippet of the raw data we use.

Table 1 shows the market shares of the five types of dark venues as a fraction of total trading volume in our sample. The total trading volume is obtained from TAQ by running the algorithm developed by Holden and Jacobsen (2014). We label the complement of these five dark venues as the “lit” venues. The “lit” label is an approximation.⁸ In the first column (full sample), we observe that dark venues account for more than 30% of total transaction volume in the 117 stocks in October 2010. Ranked by market shares, the five dark categories are *DarkRetail* (12.9%), *DarkNMid* (9.1%), *DarkOther* (6.8%), *DarkMid* (2.5%), and *DarkPrintB* (1.0%). Columns 2–4 show the breakdown of venue market shares for large, medium, and small stocks, respectively. *DarkMid* market shares are similar across the three size terciles, but *DarkNMid* market shares seem to decrease in market capitalization. *DarkRetail* market share is visibly smaller for medium stocks than for large or small ones.

It is informative to compare our five-way categorization of dark trading venues to that of the SEC. SEC (2010) classifies opaque trading centers into dark pools and broker-dealer internalization. By approximation, our *DarkMid* and *DarkNMid* types roughly fall into the SEC’s dark pool category, and our *DarkRetail*, *DarkPrintB*, and *DarkOther* types roughly fall into the SEC’s broker-dealer inter-

⁸Since NASDAQ TRF accounts for 92% of all off-exchange trading volume in our sample, the “lit” category also contains the remaining 8% of off-exchange volume, or about 2.4% of total volume.

nalization category.

Using one week of FINRA audit trail data in 2012, Tuttle (2014) reports that about 12.0% of trading volume in U.S. equities is executed in off-exchange alternative trading systems (ATSs), the majority of which are dark pools. She further reports that about 18.8% of U.S. equity volume are executed off-exchange without involving ATSs, which can be viewed as a proxy for trades intermediated by broker-dealers. The dark-pool volume and internalized volume implied by our dataset are comparable to those reported by Tuttle (2014).

Variables used in the empirical analysis. For ease of reference, all variables that are used in the empirical analysis are denoted in CamelCase italics and their descriptions are tabulated in Table 2. The sample is constructed at the minute frequency. This is necessary as the sampling frequency is constrained by NASDAQ order book snapshots, which, among all raw data used, is the one with the lowest frequency. It comprises all minutes during the trading hours (9:30am to 4:00pm) except the exact moment of market opening, 9:30am. The final dataset consists of stock-minute data with $117 \text{ stocks} \times 21 \text{ days} \times 390 \text{ minutes/day} = 958,230$ observations.

3.2 NASDAQ: trade and quote data with HFT label

Our second dataset contains detailed trade and quote data for the NASDAQ market (an electronic limit order book that is part of the “Lit” market). The data include identification of the activity of high-frequency trading firms in this market. These firms are known to engage in high-speed computerized trading, but their identities and strategies are unknown to us. This dataset has two parts.

First, for each transaction on NASDAQ, we observe the stock ticker, the transaction price, the number of shares traded, an indicator of whether the buyer or seller initiated the trade, and an indicator that for each side of the trade, tells whether it was an HFT or not. We refer to a trade as an HFT trade if at least one side of the trade is an HFT. All transactions are time-stamped to milliseconds. Second, we observe the minute by minute snapshot of the NASDAQ limit order book. For each limit order, we observe the ticker, the quantity, the price, the direction (buy or sell), and a flag on whether the order is displayed or hidden, and a flag on whether or not the order is submitted by an HFT.

An important and well-recognized caveat of the NASDAQ HFT dataset is that it excludes trades by HFTs that are routed through large, integrated brokers, nor does it distinguish various HFT strategies, be it market making or “front-running.” Given this caveat, a prudent way to interpret the HFT measures is that they are “control variables,” needed to capture market conditions. For additional details about the NASDAQ HFT dataset, see Brogaard, Hendershott, and Riordan (2014).

Table 2: Variable descriptions. This table lists and describes all variables used in this study. An underscore indicates that the variable is used in our baseline econometric model (Section 4). The subscript j indexes stocks and t indexes minutes. All volume is measured in number of shares. Type Y variables are “endogenous” variables in the econometric model (VARX) whereas type Z variables are exogenous.

Type	Variable Name	Description
Panel A: Dark venue trading volumes		
Y	$\underline{VDarkMid}_{jt}$	Dark volume in midpoint-cross dark pools
	$\underline{VDarkNMid}_{jt}$	Dark volume in non-midpoint dark pools
	$\underline{VDarkRetail}_{jt}$	Dark volume due to retail flow internalization
	$\underline{VDarkPrintB}_{jt}$	Dark volume due to average-price trades (“print back”)
	$\underline{VDarkOther}_{jt}$	Other dark volume
	\underline{VDark}_{jt}	Total dark volume as the sum of all dark volumes
	\underline{VLit}_{jt}	Total volume minus all dark volume
Panel B: NASDAQ lit market characterization		
	$\underline{ImbVolume}_{jt}$	Volume imbalance defined as the absolute difference between buyer-initiated and seller-initiated volume
	$\underline{RelSpread}_{jt}$	NASDAQ lit market bid-ask spread divided by the NBBO midpoint
	$\underline{InHiddDepth}_{jt}$	Sum of NASDAQ hidden bid depth and ask depth strictly within the displayed quotes
	$\underline{TopDispDepth}_{jt}$	Sum of NASDAQ visible best bid depth and best ask depth
	$\underline{AtHiddDepth}_{jt}$	Sum of NASDAQ hidden bid depth and best ask depth at the best quotes
	$\underline{HFTinTopDepth}_{jt}$	Percentage of depth provided by HFTs’ limit orders at or within the best quotes of the book, including hidden orders
	$\underline{HFTinVolume}_{jt}$	NASDAQ lit volume in which HFT participates divided by total NASDAQ lit volume
Panel C: Overall market conditions		
	$\underline{TAQVolume}_{jt}$	Total trading volume reported in TAQ
	$\underline{RealVar}_{jt}$	Realized variance (sum of squared one-second NBBO midquote returns)
	$\underline{VarRat10S}_{jt}$	Variance ratio: ratio of realized variance based on ten-second returns relative to realized variance based on one-second returns
Z	$\underline{dVIX^+}_{t}$	Positive innovation in VIX_t , calculated as the residual of an AR(1) estimated for ΔVIX_t
	$\underline{dVIX^-}_{t}$	Negative innovation in VIX_t , calculated as the residual of an AR(1) estimated for ΔVIX_t
	\underline{VIX}_t	VIX index level
	$\underline{Pre/PostNewsXmin}_t$	Dummy variables indicating whether an observation is within the x -minute window relative to a macro news announcement
	$\underline{PostEA1,2,\dots,13}_{jt}$	Earnings per share (EPS) announcement surprise, calculated as the absolute difference in announced EPS and the EPS forecast, scaled by share price. All announcements were done overnight. To capture the intraday response pattern the announcement variable is multiplied by a time of day dummy corresponding to the 13 half hour intervals in a trading day.

Variable used in the empirical analysis. We construct the following variables to characterize market conditions in the NASDAQ lit market. We use the absolute value of NASDAQ signed volume as a measure of volume imbalance, $ImbVolume$. Doing so with NASDAQ data has the advantage that the NASDAQ data provide the exact trade direction indicator, buy or sell. By contrast, constructing volume imbalance from TAQ transaction data would require using an imprecise trade-signing heuristic, such as the Lee and Ready (1991) algorithm. Ellis, Michaely, and O’Hara (2000) find that compared to the actual buy/sell indicator in NASDAQ data, the Lee-Ready algorithm misclassifies about 19%

of NASDAQ trades.

We further add two commonly used measures of liquidity: spread and depth. *RelSpread* is the relative bid-ask spread of the NASDAQ lit market. Three depth measures are added: *InHiddDepth*, *TopDispDepth*, and *AtHiddDepth* (ordered according to execution priority). First, *InHiddDepth* measures the hidden orders that are posted strictly within the bid-ask spread; these orders have the highest execution priority. Second, *TopDispDepth* measures the depth provided by displayed limit orders at the best bid and ask prices. Third, *AtHiddDepth* measures the hidden depth provided at the best displayed prices.

Lastly, we use two measures of HFT activity: *HFTinVolume* and *HFTinTopDepth*. *HFTinVolume* is NASDAQ volume where HFT was on at least one side of the transaction, divided by total NASDAQ volume. *HFTinTopDepth* is the number of shares posted by HFT at or within the best quotes on NASDAQ divided by the total number of shares posted at or within the best quotes on NASDAQ.

3.3 Overall market conditions

A third dataset is used to characterize overall market conditions (i.e., based on including *all* markets, not only NASDAQ). The data used for this part are the standard TAQ data and millisecond-level NBBO data provided by NASDAQ.

Variables used in the empirical analysis. For each stock and each minute, we calculate total transaction volume *TAQVolume*, realized return variance within the minute *RealVar*, and the variance ratio within the minute *VarRatioS*. The first variable is based on TAQ data, whereas the other two are based on millisecond-level NBBO data.

3.4 Proxies for urgency shocks

We proxy urgency by market-wide volatility VIX, macroeconomic data releases, and firms' earnings surprises. The first and second variable pertain to the entire market, whereas the third is firm-specific. We pick these variables for two reasons. First, VIX, macroeconomic data releases, and earnings surprises are arguably *exogenous* to the trading process of a particular stock. This is important for our purpose. Since VIX is derived from options prices and covers the entire equity market, it is unlikely to be affected by trading in an individual stock. (An individual stock's volatility, by contrast, is likely endogenous to the trading process.) Macroeconomic data releases and earnings announcements are scheduled in advance and are therefore not affected by trading activities on the announcement days.

Second, we argue that investors' urgency to trade (i.e., the opportunity cost of failing to execute orders), is higher following a surprise VIX increase, a macroeconomic data release, or an earnings

number that is far from market expectations. The motivation for using VIX is as follows. If an investor is in the market for a trade, he is exposed to risks in his current portfolio (either long or short) before the order is filled. To the extent that a higher VIX implies a general volatility increase, a risk-averse investor's cost of staying on the undesired exposure increases his utility cost of not trading. Moreover, a VIX shock, positive or negative, is by itself news that could trigger hedging trades by which securities move between relatively risk-averse investors and relatively risk-loving ones in equilibrium (see, e.g., Campbell et al., 1993). Combining the inventory-cost and volatility-news effects of VIX, we see that a positive shock to VIX unambiguously leads to a higher urgency of investors, since inventory cost and volatility news both point in this direction. But a negative shock to VIX has ambiguous implications for urgency, since inventory cost suggests lower urgency and volatility news suggests higher urgency. We therefore use positive, but not negative, VIX surprises as our proxy for urgency.

The rationale for using macroeconomic data releases and earnings announcements is similar in spirit. Following important news on the macro-economy or individual firms, investors naturally wish to adjust their positions in the stock to take into account the new information. Although macroeconomic and earnings news are made publicly, trading interests are still generated when investors close pre-news positions or interpret the same news differently (Kim and Verrecchia, 1994). Sarkar and Schwartz (2009) provide evidence that after macroeconomic and earnings news, the market generally becomes more “two-sided,” that is, the correlation between buyer-initiated trades and seller-initiated trades goes up. They conclude that this evidence supports heterogeneous opinions as a motive to trade. Under this trading motive, trading urgency goes up precisely after public news.

Data sources. VIX is a forward-looking volatility measure, calculated from option prices. The Chicago Board of Exchange (CBOE) disseminates VIX every 15 seconds. Minute by minute VIX data are obtained from pittrading.com.

Following Brogaard et al. (2014), we collect all the intraday macroeconomic data releases during October 2010 from Bloomberg.⁹ Each of these releases is accompanied by the exact scheduled announcement time (accurate to the second) and a “relevance score.” It is quite common to have multiple releases scheduled at the same time (most often at 10:00 am ET). With a selection threshold of relevance higher than 70%, we obtain 14 unique release times on 10 different trading days (out of 21) in our sample period.¹⁰ These macroeconomic data releases are used to construct the news

⁹Many macroeconomic data releases are scheduled before the market opens, say at 8:30am. As there might be surprise events between such data releases and the market open time 9:30am, we cannot be sure that the stock trading behavior at the market open time is due to the macroeconomic data release some time earlier. For this reason, we focus on macroeconomic data releases made during trading hours.

¹⁰The included macroeconomic data releases are University of Michigan Consumer Sentiment, ISM Manufacturing, ISM Prices Paid, Construction Spending, Factory Orders, Pending Home Sales, ISM Non-Manufacturing, Wholesale

announcement dummies defined below.

Of the 117 firms in our sample, 68 announced earnings in October 2010. All these announcements were made outside trading hours. For each of these earnings announcements, we download the announcement dates, time stamps, announced earnings per share, and expected earnings per share from Bloomberg. We are able to collect the EPS forecast from Bloomberg for 67 of these firms, and hence can construct the earnings surprises described below.

Variables used in the empirical analysis. From the raw data we calculate three types of exogenous variables that will be included in the empirical analysis: (i) innovations in VIX, (ii) dummies indicating minutes before and after macroeconomic data releases, and (iii) earnings surprises for stocks with earnings announcements.

As discussed above, only a positive VIX shock unambiguously leads to higher urgency. We define the minute to minute change in VIX as $dVIX(t) \equiv VIX(t) - VIX(t - 1)$. Then, we compute VIX innovations as the residuals of an AR(1) model for $dVIX(t)$ and decompose it by sign into a positive part and a negative part:

$$dVIX(t) = \alpha dVIX(t - 1) + innovation(t), \quad (1)$$

$$dVIX_+^*(t) = \max \{0, innovation(t)\}, \quad (2)$$

$$dVIX_-^*(t) = -\min \{0, innovation(t)\}. \quad (3)$$

The choice of one lag for in the autoregressive model is based on applying the BIC model selection criterion. In addition to the VIX innovation variables, we also include the level of VIX in our analysis, since both VIX changes and VIX level may affect investors' urgency to trade.¹¹

For each of the 14 unique macroeconomic data release times, we construct six time dummy variables, *PreNews1min*, *PostNews0min*, *PostNews1min*, ..., *PostNews4min*. For example, suppose there is a news release scheduled at 10:00am on October 1, which we treat as the very first event in the time interval of (10:00,10:01]. The dummy variable *PreNews1min* is set to 1 for all stocks for the minute ending at 10:00am, i.e. the minute of (9:59,10:00], on October 1, and zero otherwise; the dummy variable *PostNews0min* is set to 1 for all stocks for the minute of (10:00,10:01] on October 1, and zero otherwise; the dummy variable *PostNews1min* is set to 1 for all stocks for the minute of (10:01,10:02] on October 1, and zero otherwise; and so on.¹²

Inventories, Monthly Budget Statement, Philadelphia Fed Business Outlook, Leading Economic Index, Existing Home Sales, Dallas Fed Manufacturing, Richmond Fed Manufacturing, Consumer Confidence, FHFA House Price, New Home Sales, and Chicago Purchasing Managers.

¹¹Since VIX is a stationary process (with slow decay when sampled at the minute frequency) and it appears so during our sample period, including it in our regression model does not raise econometric concerns.

¹²Another way to construct the macroeconomic variables is to measure how far the announced numbers differ from

Consistent with the accounting literature (e.g., Kinney et al., 2002), the earnings surprise is calculated as the absolute difference between announced EPS and the pre-announcement expected EPS, divided by the closing price on the business day immediately before the announcement. Since all of our earnings announcements are made outside trading hours, it is not possible to study the immediate effect of earnings announcement on trading activity using minute-level dummies as we did with macroeconomic data releases. As such, we construct 13 intraday variables for the immediate next business day after an announcement. The dummies are $PostEA1_{jt}, \dots, PostEA13_{jt}$, one for each of the thirteen 30-minute windows in a trading day. If stock j announced its earnings on a particular day with EPS surprise of x basis points, then all of $\{PostEA1_{jt}, \dots, PostEA13_{jt}\}$ are set to x for the corresponding time window on the business day right after stock j 's earnings announcement. Otherwise, these variables are set to zero.

3.5 Data preparation and summary statistics

We convert all variables into logs, except the urgency proxies (i.e., VIX level, VIX innovation, macroeconomic news, and earnings announcements) because some of them are dummy variables. A log-linear model has a couple of advantages over a linear model. First, a log-linear model comes with a natural interpretation that estimated coefficients are elasticities. Second, all endogenous variables (e.g., volume, realized variance, and depth) are guaranteed to remain non-negative. In other words, the error term does not need to be bounded from below, which would be the case for a linear model.

To take the log, the data need to be winsorized to eliminate zeros. We use the following procedure. If a particular dark venue has zero transaction volume for stock j and minute t , its volume for that stock-minute is reset to one share. If a particular stock j does not trade in period t on the NASDAQ market, the $HFTinVolume$ variable is undefined. In this case, to avoid losing observations, $HFTinVolume$ is forward filled from the start of the day. The motivation is that market participants may learn HFT activity on NASDAQ by carefully parsing market conditions. If there is no update in a particular time interval, they might rely on the last observed value. Zero entries for all other variables are left-winsorized at the 0.01% level.

Table 3 contains summary statistics of the model variables, before taking the log. Columns 1 through 8 show the mean, standard deviation, minimum, and maximum of the raw and winsorized

market expectations. This method, however, runs into a problem if multiple macroeconomic data releases coincide at the same clock time. For example, on October 1, 2010, three macroeconomic data releases happened at 10:00am: ISM Manufacturing, ISM Prices Paid, and Construction Spending. For these situations there is no obviously optimal way to synthesize multiple releases into one “surprise” measure. For this reason, we use time dummy variables. To the extent that the surprise component is missing from these time dummies, using time dummies is conservative and goes against us finding any significant result.

data for the full sample. One important observation is that the data preparation procedure discussed above leaves the data almost unchanged as the raw data are very close to the prepared data.

The summary statistics characterize trading in our sample. Total average trading volume per stock-minute is 10,300 shares. The average bid-ask spread in the NASDAQ limit order book is 16.8 basis points. The average displayed depth at the best quotes is 4,443 shares. HFT participation in the best quotes, $HFTinTopDepth$, is 37%, whereas HFT participation in NASDAQ trades, $HFTinVolume$, is 40%. The averages of $dVIX_+^*$ and $dVIX_-^*$ are 0.7 basis points per minute, and since one of these two innovations is always zero each minute by construction, the average absolute innovation in the VIX is the sum of the two averages, or about 1.4 basis point per minute. The last six columns show the mean and standard deviation of all model variables for large, medium, and small stocks, respectively. As expected, larger stocks have higher volumes in all venues, lower spreads, higher depths, and higher HFT participation.

4 A VARX model of dark volumes

This section characterizes the dynamic interrelation of the various types of dark volume, the various measures of market conditions, and the exogenous shocks to urgency, through a panel vector autoregressive model with exogenous variables: a panel VARX.

4.1 Panel VARX model

For each stock j and each minute t , all the log-transformed endogenous variables, underscored and in the Y section of Table 2, are arranged into a vector y_{jt} . All exogenous variables, listed in the Z section of Table 2, are arranged into a vector z_{jt} . Then, the panel VARX model used for the main empirical analysis has the following form:

$$y_{jt} = \alpha_j + \Phi_1 y_{j,t-1} + \dots + \Phi_p y_{j,t-p} + \Psi z_{j,t} + \varepsilon_{jt}. \quad (4)$$

A stock fixed effect α_j ensures that only time variation is captured, not cross-sectional variation, as our focus is on dynamic interrelations among variables. We set the number of lags (in minutes) equal to two based on the BIC criterion for model selection. Further estimation details are provided in Appendix B.

4.2 Estimation results

Table 4 presents the VARX estimation results. The estimated coefficients $\{\Phi_1, \Phi_2, \Psi\}$ can be interpreted as elasticities. The results lead to a few observations. First, VIX shocks, macroeconomic

Table 3: Summary statistics. This table reports summary statistics of the variables used throughout the paper (before taking logarithm). Sample mean, standard deviation, minimum, and maximum are calculated for the full sample before and after data preparation (the first 8 columns). For each stock size tercile, subsample mean and standard deviation are then tabulated (in the other columns). The sample frequency is minute. The units of each series is in the square brackets.

	Mean (full)		StDev (full)		Min (full)		Max (full)		Mean			StDev		
	Raw	Prep/d	Raw	Prep/d	Raw	Prep/d	Raw	Prep/d	Large	Medium	Small	Large	Medium	Small
<i>VDarkMid</i> [1k shares]	0.256	0.257	2.379	2.379	0.000	0.001	1297.200	1297.200	0.709	0.045	0.017	4.005	0.730	0.315
<i>VDarkNMid</i> [1k shares]	0.938	0.939	5.411	5.411	0.000	0.001	1040.352	1040.352	2.638	0.138	0.040	9.069	1.047	0.379
<i>VDarkRetail</i> [1k shares]	1.328	1.329	6.181	6.181	0.000	0.001	1641.229	1641.229	3.783	0.133	0.070	10.219	0.894	0.574
<i>VDarkPrintB</i> [1k shares]	0.102	0.103	8.001	8.001	0.000	0.001	3524.920	3524.920	0.294	0.009	0.005	13.827	0.397	0.791
<i>VDarkOther</i> [1k shares]	0.703	0.704	10.627	10.627	0.000	0.001	3018.263	3018.263	1.948	0.108	0.057	18.136	1.926	1.942
<i>VDark</i> [1k shares]	3.437	3.438	41.120	41.120	0.000	0.001	23616.586	23616.586	9.692	0.431	0.189	70.714	2.674	2.510
<i>VLit</i> [1k shares]	7.001	7.002	31.413	31.413	0.000	0.001	9238.898	9238.898	19.306	1.272	0.427	52.102	3.781	1.972
<i>TAQVolume</i> [1k shares]	10.325	10.329	42.534	42.533	0.000	0.006	9329.493	9329.494	28.666	1.705	0.616	69.887	5.180	3.334
<i>ImbVolume</i> [1k shares]	1.282	1.282	5.943	5.943	0.000	0.001	249.800	249.800	3.476	0.274	0.096	9.883	0.936	0.417
<i>RelSpread</i> [bps]	16.776	16.776	28.466	28.466	0.162	0.334	6681.449	6681.449	4.100	12.136	34.075	2.860	12.062	42.368
<i>InHiddDepth</i> [1k shares]	0.178	0.178	1.010	1.010	0.000	0.001	249.800	249.800	0.238	0.134	0.163	1.418	0.778	0.663
<i>TopDispDepth</i> [1k shares]	4.443	4.443	13.909	13.909	0.001	0.004	502.515	502.515	11.431	1.327	0.572	22.231	3.398	1.034
<i>AtHiddDepth</i> [1k shares]	1.170	1.170	6.605	6.605	0.000	0.001	903.559	903.559	2.099	0.968	0.443	9.061	6.592	1.967
<i>HFTinTopDepth</i> [percent]	36.744	36.748	27.959	27.953	0.000	0.029	100.000	100.000	46.778	37.478	25.987	24.021	28.761	26.892
<i>HFTinVlm</i> [percent]	39.985	37.157	23.939	26.038	0.000	0.125	100.000	100.000	45.627	38.063	27.708	19.697	26.978	27.539
<i>RealVar</i> [bps]	6.139	6.188	9.112	9.079	0.000	0.172	1367.197	1367.197	5.134	5.541	8.000	4.639	7.348	12.871
<i>VarRatIOS</i> [percent]	99.772	99.778	44.438	44.426	0.000	0.317	853.407	853.407	101.613	100.834	96.905	47.554	45.698	39.431
<i>dVIX</i> *+ [percent]	0.007	0.007	0.018	0.018	0.000	0.000	0.427	0.427	0.007	0.007	0.007	0.018	0.018	0.018
<i>dVIX</i> *- [percent]	0.007	0.007	0.016	0.016	0.000	0.000	0.342	0.342	0.007	0.007	0.007	0.016	0.016	0.016
<i>VIX</i> [percent]	20.626	20.626	1.369	1.369	17.940	17.940	24.330	24.330	20.626	20.626	20.626	1.369	1.369	1.369
<i>PreNews1min</i> [1/0]	0.002	0.002	0.041	0.041	0	0	1	1	0.002	0.002	0.002	0.041	0.041	0.041
<i>PostNews0min, ..., PosteNews4min</i> [1/0]	0.002	0.002	0.041	0.041	0	0	1	1	0.002	0.002	0.002	0.041	0.041	0.041
<i>PostEA1, ..., PosteEA13</i> [percent]	0.001	0.001	0.042	0.042	0	0	4.190	4.190	0.000	0.001	0.001	0.014	0.061	0.037

data releases, and earnings surprises all forecast higher volumes in dark venues and higher total TAQ volume, but the elasticity of TAQ volume is higher than that of all types of dark volume which, in turn, show large variation. This suggests that the market shares of various venue types respond rather differently to shocks in the urgency proxies. Moreover, right after those exogenous shocks, liquidity conditions in the market tend to worsen, as indicated by a wider spread and a lower (visible) depth. Second, the various types of dark volume seem to respond very differently to changes in market condition variables such as spread, depth, and return variance, sometimes with opposite signs (as judged from the first five rows). We will return to these points later.

An informative, intuitive, and standard way to summarize the rich information in the dynamic system is to calculate and plot impulse-response functions (IRFs). In particular, IRFs reveal not only how shocks transmit across variables, but also shows how long they last. Thus, we will mostly rely on IRFs for exposition in the remainder of the paper. As the IRF is a non-linear function of parameter estimates, we establish the 95% confidence bounds of the IRF through simulations. In each iteration a value for the parameter vector is drawn from a multivariate normal with a mean equal to the point estimate and a covariance matrix equal to the estimated parameter covariance matrix. This simulation method is described in more detail in Appendix B. In the next section, we use the IRFs based on a shock to the exogenous urgency proxies to test the pecking order hypothesis. In Section 7, we present IRFs based on shocks to endogenous variables to explore heterogeneity across dark venues and across stocks.

5 Results: pecking order

In this section we test the pecking order hypothesis laid out in Section 2. We focus on testing the specific form of the pecking order hypothesis (Fig. 2(b)) because it makes stronger (i.e., more specific) predictions. Our empirical strategy is to start from the steady-state estimate of the VARX model and then study how the market shares of DarkMid, DarkNMid, and Lit respond to VIX shocks, macroeconomic data releases, and earnings surprises. In all analysis we will also discuss the behavior of large-cap, mid-cap and small-cap stocks separately.

5.1 A VIX shock

Starting with our estimated VARX model, we shock $dVIX_{\dagger}^*$ by +0.01% to examine how the volume shares of DarkMid, DarkNMid, and Lit change relative to their steady state levels.¹³ This steady state

¹³We are careful to also shock the VIX level by the amount implied by the shock to its innovation. This ensures the resulting IRF is internally consistent.

Table 4: VARX Estimation. This table presents VARX estimation results. All variables are sampled at a one minute frequency. They are log-transformed except for the exogenous variables: VIX innovation and level, earnings announcement surprises (*PostEA*) and the macroeconomic data release dummies. The (unadjusted) R^2 values are reported for each endogenous variable. Statistical inference is based on two-way clustering of the residuals, by stock and by time.

	<i>VDarkMid</i>	<i>VDarkNMid</i>	<i>VDarkRetail</i>	<i>VDarkPrintB</i>	<i>VDarkOther</i>	<i>TAQVolume</i>	<i>ImbVolume</i>	<i>RelSpread</i>	<i>InHiddDepth</i>	<i>TopDispDepth</i>	<i>AtHiddDepth</i>	<i>HFTinTopDepth</i>	<i>HFTinVlm</i>	<i>RealVar</i>	<i>VarRat10S</i>
Endogenous variables: 1-minute lag															
<i>VDarkMid</i> (-1)	0.210**	0.032**	0.026**	0.005**	0.038**	0.011**	0.010**	-0.000	0.001	0.001*	0.003	-0.000	0.001	0.002	-0.001*
<i>VDarkNMid</i> (-1)	0.028**	0.190**	0.048**	0.004**	0.045**	0.034**	0.019**	-0.000	0.004*	0.004**	0.002	0.001	0.001	0.002	-0.002**
<i>VDarkRetail</i> (-1)	0.019**	0.034**	0.110**	0.003**	0.038**	0.017**	0.021**	-0.000	-0.003*	-0.000	0.001	-0.001	0.003**	0.005**	0.001**
<i>VDarkPrintB</i> (-1)	0.010**	0.008**	0.011**	0.054**	0.019**	0.006*	0.008*	0.001	-0.002	0.001	-0.002	0.002	0.001	0.007**	0.000
<i>VDarkOther</i> (-1)	0.021**	0.036**	0.036**	0.002**	0.143**	0.017**	0.018**	-0.001*	0.001	-0.001	0.006**	-0.001	-0.002*	0.004**	-0.001
<i>TAQVolume</i> (-1)	-0.004*	0.022**	0.020**	0.001	0.007*	0.169**	0.072**	-0.002**	0.004	0.002**	0.007*	-0.008**	-0.014**	0.002	0.003*
<i>ImbVolume</i> (-1)	-0.001	-0.002	0.009**	0.001*	0.008**	0.010**	0.079**	0.000	-0.003	-0.002**	0.000	-0.003*	0.004**	0.007**	-0.001*
<i>RelSpread</i> (-1)	-0.077**	-0.120**	-0.080**	-0.003	-0.130**	-0.429**	-0.876**	0.388**	0.051*	-0.043**	-0.219**	-0.113**	-0.143**	0.189**	-0.046**
<i>InHiddDepth</i> (-1)	0.007**	0.018**	0.006**	0.001	0.018**	0.023**	0.032**	-0.002**	0.318**	0.001**	0.046**	0.004	-0.009**	-0.011**	0.000
<i>TopDispDepth</i> (-1)	0.077**	0.136**	0.071**	0.004	0.080**	0.130**	0.208**	-0.025**	-0.038**	0.372**	0.090**	-0.021	0.002	-0.108**	0.016**
<i>AtHiddDepth</i> (-1)	0.011**	0.013**	0.008**	0.001	0.015**	0.008**	0.019**	-0.002**	0.023**	0.003**	0.329**	-0.001	-0.002*	-0.011**	-0.001**
<i>HFTinTopDepth</i> (-1)	-0.003**	-0.004**	-0.002*	0.000	-0.007**	-0.001	-0.012**	-0.002**	0.001	-0.000	-0.002	0.369**	0.046**	0.004**	0.003**
<i>HFTinVlm</i> (-1)	-0.003*	-0.005**	-0.001	-0.000	-0.007**	-0.004**	-0.007**	0.000	-0.004*	-0.000	-0.000	0.004*	0.544**	-0.000	0.000
<i>RealVar</i> (-1)	-0.025**	-0.044**	-0.020**	-0.000	-0.038**	-0.045**	0.005	0.006**	-0.010**	-0.012**	-0.042**	0.011**	0.005*	0.219**	-0.023**
<i>VarRat10S</i> (-1)	0.002	-0.002	0.014**	0.002	0.001	0.006*	0.007*	0.003**	-0.000	-0.002*	-0.006	0.006*	0.005*	0.016**	0.047**
Endogenous variables: 2-minute lag															
<i>VDarkMid</i> (-2)	0.175**	0.019**	0.014**	0.003**	0.027**	0.004**	0.004	0.000	-0.000	0.002**	0.006**	-0.000	-0.002*	0.001	-0.000
<i>VDarkNMid</i> (-2)	0.016**	0.136**	0.031**	0.002*	0.029**	0.017**	0.008**	0.001**	0.007**	0.004**	0.010**	-0.000	-0.002*	-0.002	-0.000
<i>VDarkRetail</i> (-2)	0.014**	0.022**	0.088**	0.002	0.027**	0.013**	0.015**	0.000	-0.003**	0.000	-0.000	-0.000	0.001	0.006**	0.002**
<i>VDarkPrintB</i> (-2)	0.010**	0.006	0.011**	0.030**	0.007	0.001	0.003	0.000	0.004	0.001	0.007	0.001	-0.001	0.000	-0.000
<i>VDarkOther</i> (-2)	0.010**	0.021**	0.028**	0.000	0.111**	0.008**	0.010**	-0.001*	0.004**	-0.000	0.006**	-0.003*	-0.001	0.001	-0.000
<i>TAQVolume</i> (-2)	-0.012**	0.015**	0.015**	0.003**	-0.004	0.118**	0.039**	-0.004**	0.007**	0.004**	0.009**	-0.006*	-0.006**	-0.010**	0.001
<i>ImbVolume</i> (-2)	0.001	-0.002	0.008**	0.000	0.002	0.004*	0.052**	-0.004**	-0.008**	0.002**	0.009**	-0.008**	0.005**	-0.002	0.001*
<i>RelSpread</i> (-2)	0.010	0.034**	0.024**	0.002	0.027**	0.040**	0.059**	0.200**	0.005	-0.014**	-0.126**	-0.116**	-0.009	0.041**	0.009**
<i>InHiddDepth</i> (-2)	-0.000	0.002	-0.001	-0.001	0.002	0.001	-0.001	-0.001**	0.176**	0.001	0.030**	0.005**	-0.001	-0.005**	0.000
<i>TopDispDepth</i> (-2)	-0.003	0.025**	-0.005	0.003	-0.032**	-0.019**	-0.025**	-0.003	-0.009	0.217**	0.071**	-0.001	-0.009*	-0.068**	0.006**
<i>AtHiddDepth</i> (-2)	0.002	0.003*	0.001	0.000	0.002*	0.001	0.001	-0.001**	0.017**	0.002**	0.186**	0.003*	0.000	-0.009**	-0.001
<i>HFTinTopDepth</i> (-2)	0.000	-0.000	-0.001	0.000	-0.002	-0.004**	-0.008**	-0.003**	0.001	0.001	0.006*	0.192**	0.010**	-0.001	0.000
<i>HFTinVlm</i> (-2)	0.001	0.001	-0.001	-0.000	0.002	-0.000	0.004*	-0.001**	-0.003*	0.000	0.001	0.017**	0.115**	0.005**	0.001
<i>RealVar</i> (-2)	-0.010**	-0.023**	-0.010**	-0.002	-0.012**	-0.017**	0.024**	0.007**	-0.006**	-0.009**	-0.030**	0.005*	-0.001	0.128**	-0.012**
<i>VarRat10S</i> (-2)	-0.003	-0.008**	0.002	0.000	-0.005*	-0.009**	-0.008**	0.002**	-0.001	0.001	-0.006	0.002	0.003	-0.005*	0.029**

*, ** Significant, respectively, at 5%, and 1%. All tests are two sided.

Table 4 continued. . .

	<i>VDarkMid</i>	<i>VDarkNMid</i>	<i>VDarkRetail</i>	<i>VDarkPrintB</i>	<i>VDarkOther</i>	<i>TAQVolume</i>	<i>ImbVolume</i>	<i>RelSpread</i>	<i>InHiddDepth</i>	<i>TopDispDepth</i>	<i>AtHiddDepth</i>	<i>HFTinTopDepth</i>	<i>HFTinVlm</i>	<i>RealVar</i>	<i>VarRat10S</i>
Exogenous variables															
<i>dVIX*+</i>	2.337**	3.510**	3.306**	0.154	3.845**	6.874**	8.056**	0.489**	-0.108	-0.198**	-0.272	0.594**	2.129**	6.387**	0.687**
<i>dVIX*-</i>	2.704**	3.540**	2.855**	-0.172	3.777**	7.169**	9.058**	0.498**	0.075	-0.183**	-0.543	0.365	2.459**	6.956**	0.550**
<i>VIX</i>	-0.002	0.020**	-0.014*	0.007*	0.022**	0.048**	0.053**	0.002*	-0.012	0.001	0.008	-0.008	0.010**	0.045**	0.000
<i>PreNews1min</i>	-0.135*	-0.214**	-0.197**	0.118**	-0.174**	-0.282*	-0.371	0.099**	-0.024	-0.058**	-0.074	-0.485**	-0.064	-0.216	-0.032
<i>PostNews0min</i>	0.304**	0.534**	0.391**	0.021	0.532**	0.802**	0.777**	0.178**	0.251**	-0.107**	-0.071*	0.251**	0.269**	0.890**	-0.016
<i>PostNews1min</i>	-0.225**	-0.137**	-0.116*	-0.015	-0.155	-0.037	-0.010	0.035	0.100**	-0.025**	-0.070	-0.064	-0.020	-0.044	0.128**
<i>PostNews2min</i>	-0.089**	-0.031	-0.109	-0.049	-0.013	0.016	0.004	-0.063**	-0.127*	-0.009	-0.017	0.183*	0.048	-0.005	0.140**
<i>PostNews3min</i>	0.035	0.021	-0.061	0.007	0.046	0.015	0.010	-0.031**	-0.072**	-0.008	0.207**	-0.148*	-0.033*	0.098	0.085**
<i>PostNews4min</i>	-0.014	0.025	-0.021	-0.062**	0.070	-0.057	-0.054	0.025	0.027	-0.031	-0.142	0.055	-0.006	0.056	0.004
<i>PostEA1</i>	0.280*	0.315*	0.574**	0.066	0.251	0.640**	0.668**	0.008	-0.000	0.006	0.219**	-0.049	0.074	0.245*	-0.003
<i>PostEA2</i>	0.135	0.394**	0.571**	0.021	0.158*	0.551**	0.430**	0.011	0.059	-0.013	-0.022	0.104*	0.131**	0.166**	-0.045*
<i>PostEA3</i>	0.164	0.421**	0.280**	0.105	0.171*	0.479**	0.389**	0.011	0.227*	0.019	0.031	0.072	0.029*	0.171**	-0.009
<i>PostEA4</i>	0.217*	0.251*	0.238**	0.084	0.189	0.506**	0.421**	0.005	0.055	0.013	-0.029	0.092	0.090	0.236**	0.025
<i>PostEA5</i>	0.099**	0.179*	0.184**	0.080*	0.098	0.281**	0.301**	-0.018	-0.002	-0.015	0.147	0.063	0.001	0.140**	0.013
<i>PostEA6</i>	0.052	0.138*	0.169**	0.102	0.078	0.270**	0.175	-0.023	-0.059	0.000	0.070	-0.026	-0.003	0.062	0.017
<i>PostEA7</i>	0.048	0.129*	0.244**	0.077*	0.131*	0.273**	0.214*	0.018	0.246**	0.009	-0.034	0.184**	0.016	0.053	0.015
<i>PostEA8</i>	0.105**	0.242*	0.252**	0.088*	0.160*	0.252*	0.254*	0.019	0.163**	0.015	-0.105	0.115	0.048	0.192**	0.003
<i>PostEA9</i>	0.066*	0.154**	0.133	0.100*	0.199**	0.244**	0.235*	0.001	-0.077*	0.001	-0.104	-0.068	-0.027	0.052	-0.002
<i>PostEA10</i>	0.068*	0.176*	0.217*	0.102*	0.159*	0.213*	0.188	0.008	0.036	0.000	0.001	0.055	0.113*	0.080	0.024
<i>PostEA11</i>	0.062	0.202*	0.121	0.111*	0.017	0.295**	0.195	0.007	0.017	-0.011	0.030	0.025	0.095**	0.061	-0.003
<i>PostEA12</i>	0.032	0.128	0.136	0.100	0.079	0.291**	0.166*	0.018	0.057	0.010	0.045	0.021	0.026	0.084	0.003
<i>PostEA13</i>	0.067	0.213	0.232*	0.076*	0.154	0.296**	0.192*	0.009	0.017	0.008	0.155	0.012	0.033	0.061	0.060
<i>R</i> ²	0.114	0.111	0.055	0.005	0.077	0.107	0.071	0.277	0.184	0.286	0.215	0.236	0.398	0.122	0.008
# obs.	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563	943563

*, **, Significant, respectively, at 5%, and 1%. All tests are two sided.

level is equal to the overall average volume share. While the VARX model is specified in terms of volume, the calculation of market shares is straightforward from the estimated coefficients.¹⁴ A shock size of +0.01% is approximately equal to a one standard deviation of $dVIX_+$. The volume shares are denoted by *SDarkMid*, *SDarkNMid*, and *SLit*, respectively. The pecking order hypothesis stated in Section 2 predicts that, after VIX shocks, the proportional changes of Lit, DarkNMid, and DarkMid market shares are positive, mildly negative, and most negative, respectively. The model developed in Section 8 develops this prediction formally.

The left-most column in Fig. 3 depicts our the findings for the overall sample. The ordering of the three venue types conforms to the pecking order hypothesis. In the contemporaneous minute of the VIX shock, *SDarkMid* shows the most negative reaction, falling from a steady state level of 2.40% to 2.29%, a 4.6% reduction. *SDarkNMid* also falls from a steady state level of 7.52% to 7.27%, but the fractional loss is smaller: 3.3%. By contrast, *SLit* increases from 76.2% to 77.0%, a gain of 1.0%. The reduction in market shares of DarkMid and DarkNMid remains significant over the two/three-minute horizon, but the response of Lit is significant only in the minute of the VIX shock. For all three venue types, the effects on market shares dies out completely within five minutes.¹⁵

To examine which stocks “drive the result” we estimate the VARX model separately for large, medium, and small stocks, and for each size tercile we repeat the analysis of shocking VIX by +0.01%. The three right-most columns in Fig. 3 depict our findings. We see that for medium and small stocks the pecking order hypothesis is generally supported: DarkMid and DarkNMid lose market share and Lit gains market share, and the magnitudes of market share changes relative to the steady state are comparable to or larger than their full-sample counterparts. By contrast, for large stocks the market shares of the three venues do not respond to VIX shocks in a statistically significant manner. As we see will shortly, this lack of significance for large stocks also applies to macroeconomic data releases and earnings announcements. Possible explanations are discussed at the end of this section.

The sorting of the three venue types shown in Fig. 3 is supported by formal econometric tests of the following two null hypotheses:

Null 1: The proportional changes of *SDarkMid* and *SDarkNMid* to VIX shocks are the same;

Null 2: The proportional changes of *SDarkNMid* and *SLit* to VIX shocks are the same.

We perform the tests by constructing the 95% confidence bounds on the differences between the

¹⁴Specifically, we use the estimated VARX coefficients and their covariance matrix to simulate 10,000 processes. In each simulation, we calculate the contemporaneous and subsequent volume changes following a VIX shock at the steady state. These volume changes then translate into shocked market share series in each simulation. We then compute, and plot in Fig. 3 below, the mean and the confidence bounds of the across the 10,000 simulations.

¹⁵As discussed in Section 2, the pecking order hypothesis does not make unambiguous predictions for a negative VIX shock. In unreported results, we find that a negative VIX shock leads to a smaller dark market share and a larger lit market share, although the statistical significance is borderline and the economic magnitude is smaller.

proportional changes in market shares. These confidence bounds are plotted in Fig. 4. In the full sample, after a +0.01% shock in $dVIX^*$, the percentage change of $SDarkMid$ is significantly more negative than that of $SDarkNMid$ in the contemporaneous minute and in the first two minutes, thus rejecting Null 1. The percentage change of $SDarkNMid$ is significantly more negative than that of $SLit$ contemporaneously and only marginally significant in the first minute, thus rejecting Null 2. For large stocks, Null 1 is rejected also for minute 1 to 3 (marginally so at minute 2), but Null 2 is not. For medium and small stocks, the rejection of Null 1 is slightly weaker (after the contemporaneous minute), while for Null 2 the rejection is very strong. Overall, the evidence from VIX shocks supports the pecking order hypothesis, except the DarkNMid-Lit step for large stocks.

Figure 3: Pecking order following a shock in VIX. This figure plots the impulse response functions of the market shares of DarkMid, DarkNMid, and Lit following a +0.01% shock to $dVIX^*$. The 95% confidence bounds are constructed by simulation. Panel (a) shows the result for all stocks and panels (b) through (d) show the results by stock size category. Minute 0 after the shock corresponds to the contemporaneous minute.





Figure 4: Testing the pecking order hypothesis: VIX shock. This figure plots the difference between the percentage market share changes of DarkMid and DarkNMid following a +0.01% VIX shock in the left column and the difference between DarkNMid and Lit in the right column. In each panel, the point estimate of the difference (solid line) and the 95% confidence bounds (shaded area) are shown for all stocks and separately for all stock size categories. The confidence bounds are constructed by simulation. Minute 0 after the shock corresponds to the contemporaneous minute.



5.2 Macroeconomic data releases

Recall that the effects of macroeconomic data releases on dark market shares are captured by the six time dummy variables in the VARX. The left most column in Fig. 5 plots the market shares of DarkMid, DarkNMid and Lit in the six time windows, estimated using the full sample. In the contemporaneous minute of macroeconomic data releases, DarkMid and DarkNMid market shares see distinctive drops of 38.5% and 22.7% relative to their steady state levels, whereas the Lit market share increases by 9.1%. In all three venues, the 95% bounds show that the first minute response is statistically significant for all market shares. The response lasts for about two to three minutes before becoming indistinguishable from the steady state.

We repeat the analysis for the three stock size terciles, and the results are shown in the three right most columns of Fig. 5. Just like the VIX result, medium and small stocks show strong support for the pecking order hypothesis, whereas market share changes are insignificant for large stocks. For medium stocks and in the minute after macroeconomic news, DarkMid and DarkNMid market shares are at 40% and 66% of their steady state levels, respectively. For small stocks and in the minute after macroeconomic news, DarkMid and DarkNMid market shares are 44% and 58% of their steady state levels.

As in the VIX analysis, we also formally test the two null hypotheses: after macroeconomic data releases, (i) the proportional changes of $SDarkMid$ and $SDarkNMid$ are equal, and (ii) the proportional changes of $SDarkNMid$ and $SLit$ are equal. The simulated confidence bounds on the differences in proportional market share changes are plotted in Fig. 6. The two nulls are rejected in the full sample and for medium stocks. Null 2 is also rejected for small stocks, albeit only in the first minute after the news release. Summarizing, the evidence from macroeconomic data releases also supports the pecking order hypothesis, except for large stocks.

Figure 5: Pecking order around macroeconomic data releases. This figure plots the venue market shares from one minutes before to four minutes after macroeconomic data releases. The 95% confidence bounds are constructed by simulation. Panel (a) shows the results all stocks and panels (b) through (d) show them by stock size category. Minute 0 after the shock corresponds to the contemporaneous minute.





Figure 6: Testing pecking order hypothesis: Macro news releases. This figure plots the difference between the percentage market share changes of DarkMid and DarkNMid around macro data releases in the left column and the difference between DarkNMid and Lit in the right column. In each column, the point estimate of the difference (solid line) and the simulated 95% confidence bounds (shaded area) are shown separately for the full sample and the three stock size size categories. The horizontal axis labels indicate the minute relative to the announcement time. Minute 0 after the shock corresponds to the contemporaneous minute.



5.3 Earnings announcements

The third shock we explore is the earnings surprises of individual firms. Starting with the estimated VARX model, we shock the earnings surprise by 1% (similar in size to the EPS surprise standard deviation: 0.89%) and calculate the new steady state market shares of DarkMid, DarkNMid, and Lit. Since 13 intraday effect variables are included in the VARX model, we are able to identify an intraday pattern of how the firm-specific urgency proxy affects market share.

The left most column of Fig. 7 plots the results for the full sample. A 1% higher earnings surprise significantly reduces $SDarkMid$ by about 20%, on average, throughout the day. A similar pattern but of smaller magnitude is seen for $SDarkNMid$, while $SLit$ shows overall increases. For both $SDarkMid$ and $SDarkNMid$, the effects are stronger in the earlier trading hours. While most intraday effects are not statistically significant (except for some $SDarkMid$ estimates), the point estimates echo what we found for VIX shocks (Fig. 3) and macro data releases (Fig. 5). The limited statistical significance here is due to low power of the test as we only have 67 such stock-day observations (out of 117×21). The hypothesis tests similar to Fig. 4 and 6 do not show any statistical significance and are therefore not included.

The right most columns of Fig. 7 plot the results for the large, medium, and small stocks, respectively. Consistent with the patterns around VIX shocks and macroeconomic data releases, DarkMid shows a significant drop in its market share after surprise earnings announcement for medium and small stocks, but not for large stocks. Since the number of firms that have earnings announcements is significantly reduced by partitioning the full sample into size terciles, these results are more telling about how strong the statistical significance is for medium and small stocks rather than how weak it is for large stocks.


5.4 A discussion of large stocks

A consistent empirical pattern observed in the three tests above is that the pecking order hypothesis is supported by medium and small stocks, but venue shares of large stocks are statistically unresponsive to these exogenous shocks. In this section we propose two candidate explanations.

The first candidate explanation is that because large stocks on average have lower betas than medium and small stocks, large stocks should respond less to market-wide urgency shocks like a higher VIX or macroeconomic data releases. Indeed, the average betas of the subsamples are decreasing in stock sizes: 1.37 for large, 1.12 for medium, and 1.01 for small stocks.¹⁶ We find further evidence in support of this conjecture by repeating the same analysis for the beta-sorted terciles. Af-

¹⁶The betas are calculated from a one-factor CAPM model, using the daily stock returns in the one-year period before our sample month (October 2010). Since we only have 117 stocks, the average beta of these stocks need not be 1.

Figure 7: Pecking order following a shock in earnings surprises. This figure plots the steady-state values of venue market shares for days with no earnings announcements (EA) and the corresponding venue market shares, for all 13 half-hour trading intervals, after a 1% EPS surprise. The 95% confidence bounds (2.5% and 97.5%) are constructed by simulation. Panel (a) shows the result for the full sample, while separates results for small, medium, and large stocks are shown in panels (b), (c), and (d).



ter an upward VIX shock, DarkMid loses market share while Lit gains market shares in all three beta-sorted terciles; both effects are statistically significant. While the point estimates for DarkNMid suggest that it also loses market share after the VIX shock, only in the high-beta tercile is the effect statistically significant. This set of results is suppressed to conserve space but are available upon request. Given this evidence, it appears that the lack of responsiveness of large stocks to urgency shocks can be partly, but not entirely, attributed to their low systematic risk (low beta).


The second candidate explanation is that large stocks are so liquid and traded so fast that effective delay cost in dark venues caused by execution risk is very small. Some formal support for this conjecture is derived in Section 8, which analyzes a simple model to derive the pecking order hypothesis in equilibrium. The analysis shows that the pecking order pattern is weaker if the spread is lower or volume is higher. To examine this conjectured weaker result for liquid stocks, we sort the 117 stocks by trading volume (measured in September 2010), and the three volume terciles are very close to the corresponding size terciles. As expected, the results for the volume terciles are the same as those for size terciles: evidence from medium-volume and low-volume stocks supports the pecking order hypothesis, but venue shares for high-volume stocks show no significant response. We repeat the same exercise using relative-spread sorted terciles and absolute-spread sorted terciles, and the results are the same: urgency shocks affect the venue shares of the high- and medium-spread terciles, but not the low-spread tercile. The volume sorted and spread sorted results are not reported to conserve space but are available upon request. This evidence suggests that the superior liquidity of large stocks could be an explanation why their venue shares are less sensitive to urgency shocks.

6 A brief discussion of DarkRetail

So far in the paper we have focused on DarkMid and DarkNMid. They are the key ingredients of the pecking order hypothesis and the focus of recent theories of dark pools. In this section, we briefly discuss how DarkRetail responds to the same three shocks. DarkRetail is the largest single category of dark venues, and the behavior of retail internalization is of key interest to regulators.

Fig. 8 shows how the DarkRetail market share responds to a 0.01% VIX shock (row 1), to a macroeconomic data release (row 2), and to a 1% earnings surprise (row 3). The four columns correspond to the full sample, large stocks, medium stocks, and small stocks, respectively. Three features stand out. First, the market share of DarkRetail drops after VIX shocks and macroeconomic data releases. Second, these results are largely driven by medium and small stocks; large stocks are unresponsive. These two patterns are very similar to those observed for DarkMid and DarkNMid. The third feature of DarkRetail is that there is no statistically significant reduction in DarkRetail market share in any of the size terciles following a surprise in earnings announcement, although the point

Figure 8: Impulse response of DarkRetail venue share. This figure depicts the impulse response function for DarkRetail market share following shocks to the three exogenous variables: a 0.01% VIX shock, a macro data release, and a 1% earnings announcement surprise (corresponding to rows 1, 2, and 3, respectively). The four columns correspond to the full sample, large, medium, and small stocks, respectively. Minute 0 after the shock corresponds to the contemporaneous minute.



estimates suggest that both medium and small stocks see sizable market share drops throughout of the trading day. The statistical insignificance here may be a consequence of a small sample of firm-days with earnings announcements. Statistical significance aside, it may also suggest that retail investors tend to trade more after firm-specific news than after VIX shocks or macroeconomic news, relative to institutional investors.

7 Dark market shares and market conditions

We have shown that DarkMid, DarkNMid, and DarkRetail have reduced market shares after exogenous urgency shocks. The pecking order hypothesis of DarkMid, DarkNMid and Lit is supported in the data. These results are primarily driven by medium and small stocks.

In this section we turn to the relation between dark venue shares and *endogenous* market conditions, such as spread, depth, volume, and volatility. We study how shocks to market conditions predict dark pool market shares in subsequent minutes. Because market conditions and dark pool activities are *endogenous*, we will not draw any causal conclusions. That said, documenting the dynamic relation is still a valid and interesting exercise. As before, we focus on DarkMid, DarkNMid, and DarkRetail. All impulse-response functions are calculated by starting the VARX at the steady state and shocking one variable at a time. This analysis is run for the full sample of stocks and for three size terciles separately.

Fig. 9 shows market shares of DarkMid, DarkNMid, and DarkRetail respond to an upward shock in relative spread on NASDAQ, displayed depth at the best quote on NASDAQ, total TAQ volume, and realized return variance. The shock sizes are set to the mean level of the respective variables. Since these variables are in their logarithms in the VARX model, the shock size is $\ln 2$. Such shock sizes are not particularly large given the standard deviations of these variables (see Table 3). Each panel corresponds to a shock in one of the market conditions variables. In each panel, the three rows correspond to the three dark venue types, and the four columns correspond to the four possible sample selections.

An inspection of the 48 (12×4) subplots reveals the following patterns. First, there is heterogeneity across stock size regarding how dark market shares respond to shocks in market conditions. Second, within the three dark venues, the responses of DarkMid and DarkNMid seem mostly consistent with each other but from time to time differ from that of DarkRetail. We discuss these observations in detail below and provide tentative explanations that are motivated from the existing literature.

Observation 1: Heterogeneity across large, medium, and small stocks.

- For medium and small stocks, DarkMid and DarkNMid generally gain market share following

low liquidity conditions, such as a high Lit spread, low Lit depth, low total volume, and high return variance. The only exceptions are medium stocks in DarkNMid after shocks in *TAQVolume* and *RealVar*.

- For large stocks, DarkMid and DarkNMid generally gain market share following high liquidity conditions, such as low Lit spread, high Lit volume, high total volume, and low return variance.

Clearly, large stocks behave very differently from medium and small stocks. We have seen heterogeneity across the size terciles in their response to exogenous shocks in Section 5. To explain this heterogeneity in economic terms, we like to highlight one feature from recent dark pool theories which is the degree to which investors seeking liquidity post limit orders in lit venues. In the models of Hendershott and Mendelson (2000), Degryse, Van Achter, and Wuyts (2009), and Zhu (2014), the lit venue is operated by market makers who provide quotes, and investors either send market orders or use the (midpoint) dark pool. In these models, a wider spread encourages (uninformed) investors to use the dark pool more frequently.¹⁷ The evidence from medium and small stocks supports these models. By contrast, investors in the model of Buti, Rindi, and Werner (2015) can post limit orders in the lit venue, in addition to sending market orders and using the dark pool. This implies that when the order book has a wide spread or a low depth, incoming investors prefer posting limit orders to earn the spread, which lowers dark pool usage. The evidence from large stocks supports their model.

Viewed this way, the heterogeneity across size terciles has an intuitive interpretation. For investors who are not financial intermediaries, posting limit orders is a more attractive execution strategy for large stocks because they trade highly frequently, so the model of Buti, Rindi, and Werner (2015) is more applicable. But using limit orders for relatively inactive medium and small stocks may involve a long waiting time, so the investors are better off as “liquidity takers.” For those stocks the models of Hendershott and Mendelson (2000), Degryse, Van Achter, and Wuyts (2009), and Zhu (2014) are more applicable. This interpretation parsimoniously rationalizes the first two rows of each panel of Fig. 9.

As a side point, our time-series results on dark pool market shares should be distinguished from the cross-section evidence in the prior literature, such as Buti, Rindi, and Werner (2011) and Ready (2014). Using voluntarily reported data in 11 U.S. dark pools, Buti, Rindi, and Werner (2011) find that dark pool market shares are higher for stocks with lower spreads, higher depths, higher trading volume, and lower volatilities. Ready (2014) finds that Liquidnet and POSIT are used more for stocks with less adverse selection (a component of stock volatility driven by order flows) and lower percentage spread. Together, the cross-section evidence suggests that dark pools tend to have higher

¹⁷The informed strategies are a little involved. Zhu (2014) shows that if adverse selection is severe, a smaller fraction of informed investors use the dark pool because their one-sided orders imply a low execution probability.

market shares for liquid stocks. Therefore, the liquidity-dark pool market share relation in the cross section and in the time series agree on large stocks but disagree on medium and small stocks.

Observation 2: DarkRetail.

- DarkRetail gains market share following low liquidity conditions, such as high Lit spread, low Lit depth, and high return variance. The exception is that retail activity also goes up following high volume and high depth in large stocks.

The DarkRetail results in Fig. 9 are intuitive. DarkRetail reactions to spread, depth, and variance are similar to dark pool reactions to the same variables. For example, a high spread is followed by a higher DarkRetail market share. Intuitively, since retail orders are uninformed on average and dealers handling these orders only give retail investors a small price improvement, a wider spread increases the profit of absorbing retail order flows. A similar interpretation holds for the DarkRetail reaction to shallower depth. Interestingly, retail investors seem to become more active following high individual stock volatility for medium and small stocks or high trading volume or depth for large stocks. These patterns are consistent with Barber and Odean (2007)'s finding that individual investors are particularly active in attention-grabbing stocks, such as stocks in the news, stocks experiencing high abnormal trading volume, and stocks with extreme one-day returns.

Figure 9: Response of venue shares to market condition shocks. This figure shows the impulse response functions of three venue shares, DarkMid, DarkNMid, and DarkRetail, to four types of shocks: *RelSpread* in panel (a), *TopDispDepth* in panel (b), TAQVolume in panel (c), and *RealVar* in panel (d). In each panel, twelve figures are shown in a three by four matrix, where each column indicates a different sample and each row indicates the three venues: DarkMid, DarkNMid, and DarkRetail. The shock sizes are set to the sample mean of these variables (see Table 3).

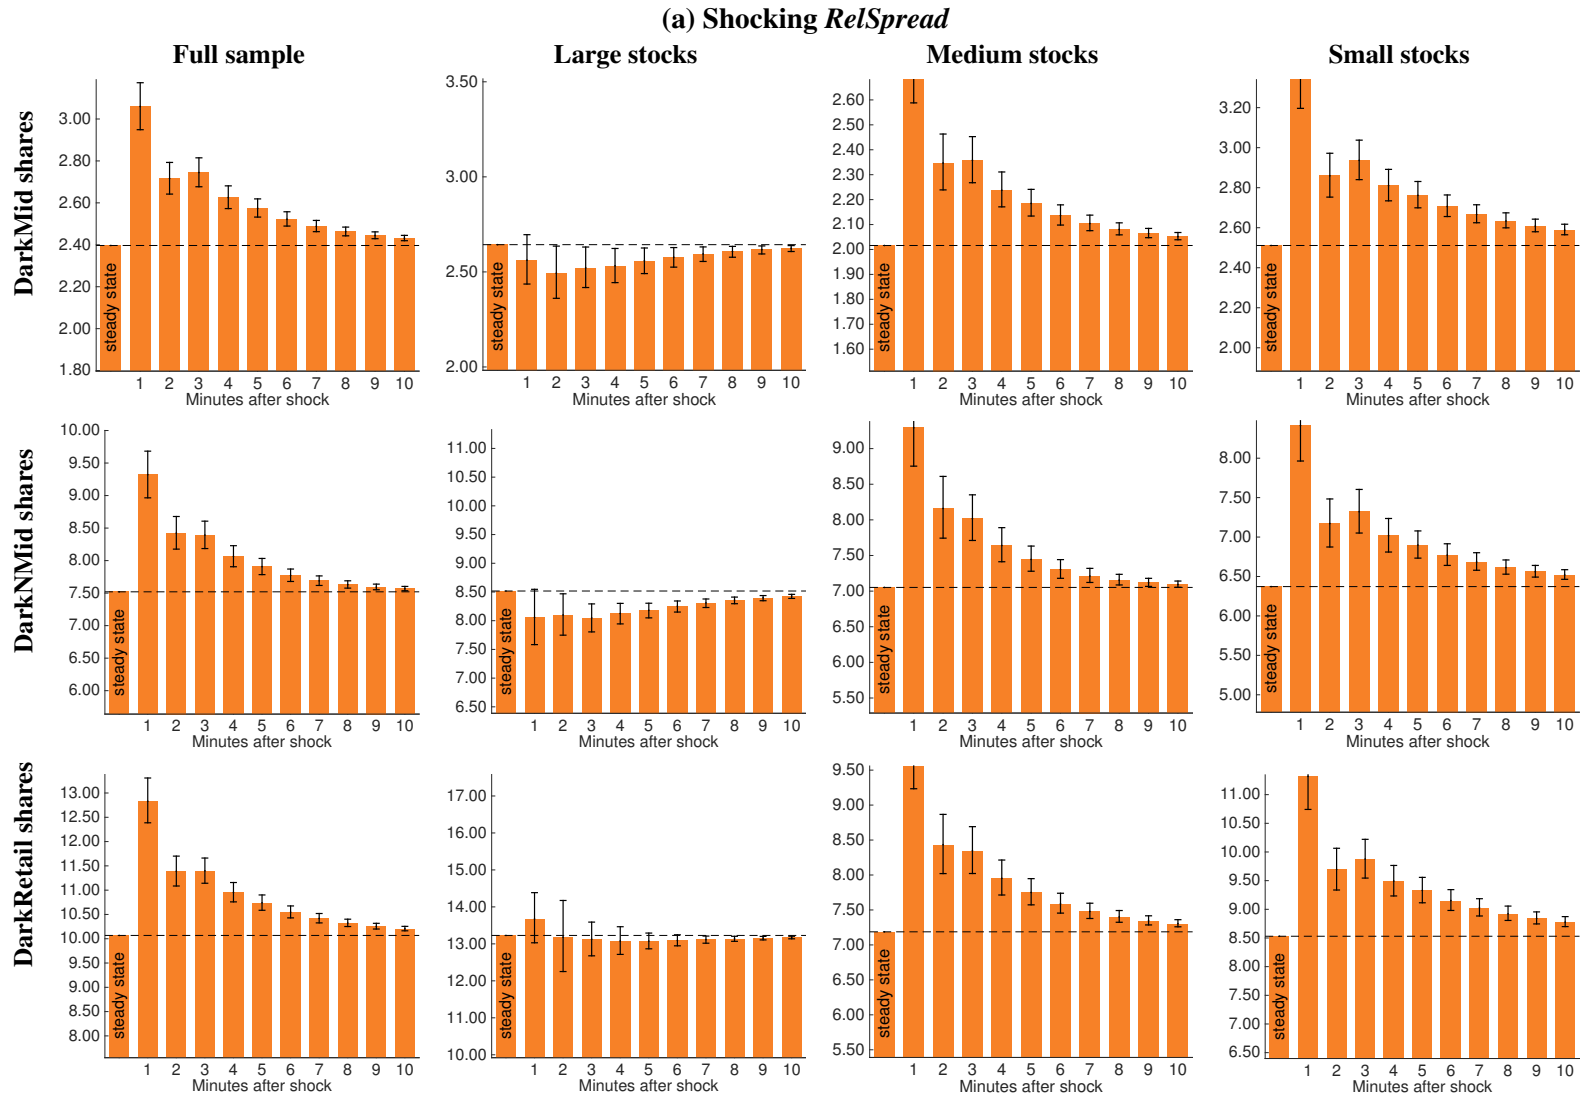


Figure 9 continued ...

(b) Shocking *TopDispDepth*

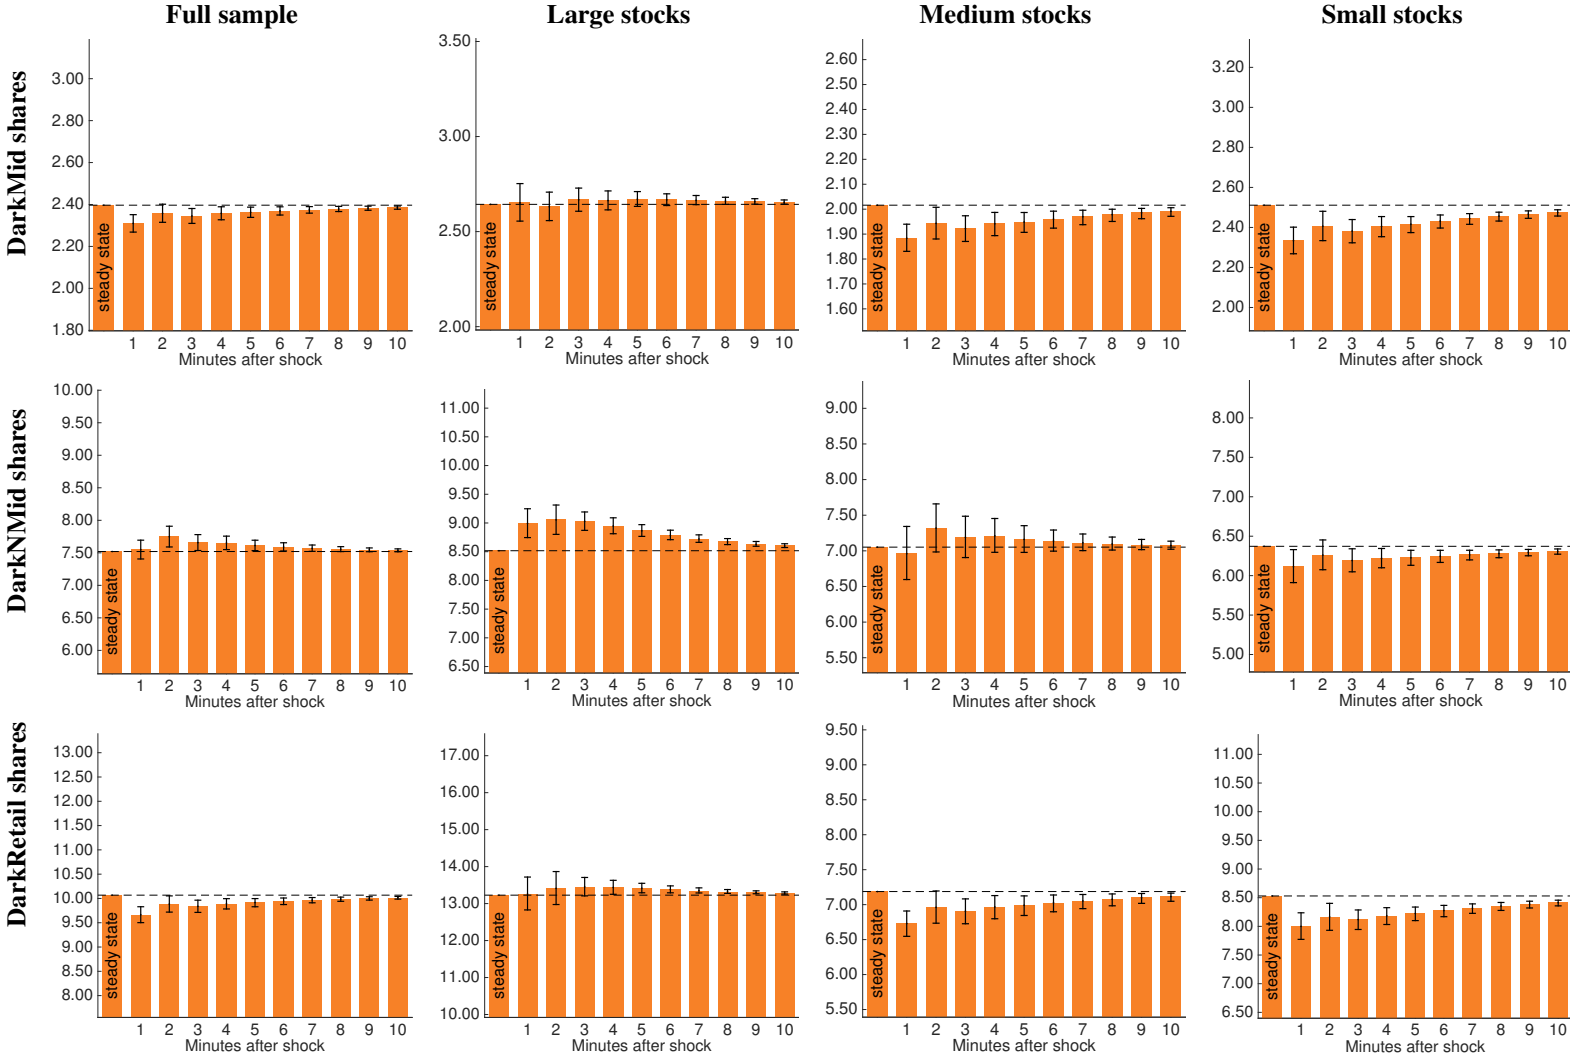


Figure 9 continued ...

(c) Shocking *TAQVolume*





Figure 9 continued ...

(d) Shocking *RealVar*



8 Pecking order as an equilibrium outcome

This section shows that the pecking order hypothesis obtains in equilibrium in a simple stylized model. Section 8.1 introduces the model and develops the result. Section 8.2 analyzes the benchmark case in which all venues are consolidated into a single one. Section 8.3 then compares the fragmented market structure and the consolidated one in terms of investor transaction cost and discusses the difference.

8.1 A simple stylized model of fragmented market trading

This section proposes a simple model that characterizes investors' choices among three venue types: DarkMid, DarkNMid, and Lit. Relative to existing theories of dark pools, our simple model distinguishes different types of dark venues. The model and its analysis formalize the intuition that led to the pecking order hypothesis.

8.1.1 Model setup

Asset. There is one traded asset. Its fundamental (common) value is normalized to zero. All players in this model have symmetric information about the asset and value it at zero. To formalize a pecking order hypothesis based on the urgency of trades, a symmetric-information setting suffices.

Investors and timing. There are potentially two investors, a buyer who has an inventory shock $-Q < 0$ units of the asset, and a seller who has an inventory shock $Q > 0$ units. While the liquidity-driven trading demand Q is common knowledge, the presence of the buyer or seller is uncertain. Specifically, a buyer is present with probability $\phi \in (0, 1)$, and likewise for the seller. The presence of the buyer is independent of the presence of the seller. Thus, conditional on the presence of one, the probability of the presence of the other is also ϕ .

There is one trading round. Before trading, Nature determines whether the buyer is present and whether the seller is present. If present, an investor then chooses the optimal trading strategies in the three venues explained below. Venues execute trades simultaneously according to their specific trading protocols explained below.

If an investor (buyer or seller) is left with a non-zero inventory q after the trading round, he incurs a quadratic cost of $(\gamma/2)q^2$, with $\gamma > 0$. Here, γ could be an inventory cost, a proxy for risk-aversion, or the cost of a missed opportunity to trade on a short-lived private signal. The parameter γ is the key parameter of the model. We interpret it broadly as investors' urgency to trade: The higher is γ , the larger is the cost of holding a non-zero net position, and hence investors are more eager to trade.¹⁸

¹⁸For example, the interpretation of risk aversion is consistent with Campbell et al. (1993), who argue that the market's

Venues and trading protocols. There are three trading venues: Lit, DarkMid, and DarkNMid.

- Lit is populated by infinitely many competitive and infinitesimal liquidity providers who have the same marginal cost β (> 0) for taking on one unit of the asset per capita, either long or short. The cost can be an operation cost or a margin cost. Together, these liquidity providers supply infinite depth at prices β and $-\beta$. This construct is similar to the “trading crowd” assumption in, for example, Seppi (1997) and Parlour (1998).

If present, the buyer’s (seller’s) strategy in Lit is represented by the size of the market buy order x_L^+ (size of the market sell order x_L^-).

- DarkMid crosses buy and sell orders at the midpoint price, i.e., at 0. If unbalanced, only the matched part of the order flow gets executed. For example, if there are buy orders for 100 units in total and sell orders for 40 units then only 40 units are matched and executed.

If present, the buyer’s (seller’s) strategy in DarkMid is represented by the size of the buy order x_M^+ (size of the sell order x_M^-).

- DarkNMid is run by a single competitive liquidity provider who starts with inventory zero, but incurs an inventory cost of $-\eta y^2/2$ for taking a net (long or short) position of y , where $\eta > 0$. The liquidity provider quotes price p_b and quantity q_b for the buyer, and price p_s and quantity q_s for the seller.¹⁹

If present, the buyer’s (seller’s) strategy in DarkNMid is represented by the size of his buy order x_N^+ at price p_b (size of his sell order x_N^- at price p_s).

A brief remark on model assumptions. This stylized model omits asymmetric information of the asset value. This restrictive assumption is made for tractability. Zhu (2014) models the coexistence of a midpoint dark pool and an exchange, allowing information asymmetry. He finds that as long as informed traders use the dark pool, the dark pool market share tends to decrease in the value of proprietary information. Solving a model that combines asymmetric information with the three types of venues is hard and beyond the empirical focus of this paper.

8.1.2 Equilibrium

Without loss of generality, we fix the seller’s strategy x_i^- and consider the buyer’s choice of x_i^+ when the buyer is present. Note that from the buyer’s perspective, the seller’s order sizes x_i^- are Bernoulli

aggregate risk aversion—which is correlated with VIX, for example—reflects the change of (a subset of) individual investors’ risk aversion.

¹⁹Our main results on the pecking order hypothesis are robust to various ways of modeling DarkNMid. For instance, the DarkNMid provider can also post two price schedules.

random variables: The seller is present only with probability ϕ . Let $V_M^+ := \min\{x_M^+, x_M^-\}$ be the trading volume in DarkMid. Then, the buyer's expected profit is

$$\pi^+ = \underbrace{-\mathbb{E}\left[0 \cdot V_M^+\right]}_{\text{Price to pay in DarkMid}} \underbrace{-p_N^+ x_N^+}_{\text{Price to pay in DarkNMid}} \underbrace{-\beta \cdot x_L^+}_{\text{Price to pay in Lit}} \quad (5)$$

$$+ \underbrace{\mathbb{E}\left[0 \cdot (Q - V_M^+ - x_N^+ - x_L^+)\right]}_{\text{Liquidation value of remaining position}} - \underbrace{\frac{\gamma}{2} \mathbb{E}\left(Q - V_M^+ - x_N^+ - x_L^+\right)^2}_{\text{Quadratic cost for failing to trade}}, \quad (6)$$

which can be simplified to

$$\pi^+ = -\frac{\gamma}{2} \mathbb{E}\left(Q - V_M^+ - x_N^+ - x_L^+\right)^2 - p_N^+ x_N^+ - \beta x_L^+. \quad (7)$$

The buyer thus chooses three parameters, x_M^+ , x_N^+ , and x_L^+ , to maximize his expected profit (7).

We will focus on a symmetric-strategy equilibrium, i.e., the buyer or the seller, if present, chooses the same order sizes (but different signs): $x_i^+ = x_i^-$ for all venues $i \in \{M, N, L\}$. And the DarkNMid provider chooses prices $p_b = p_N > 0$ and $p_s = -p_N$, and quantities $q_b = q_s$. In the equilibrium we characterize, the DarkNMid provider's offered quantities are greater than or equal to the investors' DarkNMid order size, i.e., $q_b = q_s \geq x_N$.

Because we look for a symmetric equilibrium, from this point on we suppress the superscript “+” or “-” unless we need to explicitly distinguish a buyer from a seller.

Proposition 1 (Equilibrium order flows). *If*

$$Q < \Delta := \frac{\beta}{1 - \phi} \left(\frac{1}{\gamma} + \frac{1}{\gamma + \eta} \right) \quad (8)$$

then there exists an equilibrium with the following strategies:

$$x_M = \frac{\gamma + \eta}{2\gamma + \eta} Q, \quad x_N = \frac{\gamma}{2\gamma + \eta} Q, \quad x_L = 0, \quad p_N = (1 - \phi)\gamma Q \frac{\gamma + \eta}{2\gamma + \eta}, \quad (9)$$

and q_b and q_s are arbitrary quantities not smaller than x_N .

If $Q \geq \Delta$, then there exists an equilibrium with the following strategies:

$$x_M = \frac{\beta}{(1 - \phi)\gamma}, \quad x_N = \frac{\beta}{(1 - \phi)(\gamma + \eta)}, \quad x_L = Q - x_M - x_N, \quad p_N = \beta, \quad (10)$$

and q_b and q_s are arbitrary quantities not smaller than x_N .

Let us now briefly discuss the intuition of the equilibrium. If an investor is present, for example a

buyer, he faces various tradeoffs when minimizing his expected trading cost for the liquidity shock Q . The following heuristic argument illustrates how he balances various costs across the venues.

The buyer's expected marginal cost of sending in an order to DarkMid is $(1 - \phi)\gamma x_M$: With probability ϕ , his order can be matched with the seller; and with probability $(1 - \phi)$, there is no match and he suffers inventory cost from the remaining inventory. The marginal cost of sending in an order to DarkNMid is p_N . In equilibrium, we should have $(1 - \phi)\gamma x_M = p_N$. The marginal cost of sending in an order to Lit is β .

For relatively small Q , we show in the proof that $(1 - \phi)\gamma x_M = p_N < \beta$, which implies that the buyer will not use Lit and, therefore, $x_N = Q - x_M$. Given this demand, the DarkNMid provider solves the optimal p_N and q_N to maximize profit. This case corresponds to $Q < \Delta$ in Proposition 1.

When Q is large enough, we have a corner solution of $(1 - \phi)\gamma x_M = p_N = \beta$. The buyer then finds the optimal x_N at $p_N = \beta$, knowing that he could trade the rest of his position $Q - x_M - x_N$ at the same price β in Lit.

Note that in both cases, the order sizes sent to the three venues add up to Q , and there is no reason to send more than Q .²⁰

Finally, we note that the equilibrium of Proposition 1 is not unique. For example, there exists another, less interesting equilibrium in which neither the buyer nor the seller uses DarkMid, which has the counterfactual implication that DarkMid has zero market share.

8.1.3 Urgency elasticity of venue market shares

To establish the pecking order hypothesis, we are interested in how the market share of each of the three venues responds to a change in investor urgency. We first compute the expected market shares in equilibrium and then rank the venues according to their market share elasticities with respect to the urgency parameter, γ .

The focus on market shares as opposed to raw volume is consistent with existing empirical studies of dark pools and fragmentation (see, for example, O'Hara and Ye, 2011; Buti et al., 2011; and Ready, 2014) as well as the empirical tests conducted in Section 5.

A positive γ shock is naturally thought of as the model equivalent of VIX shocks, macroeconomic

²⁰This is because an excess order sent to DarkMid is unfilled for sure, and an excess order sent to DarkMid and DarkNMid incurs positive costs. Although an investor either executes the entire desired quantity Q or strictly less, he does balance the risk of trading "too much" and trading "too little." Conditional on both investors being present, both apparently send "too little" to DarkMid—"too little" because, ex post, they could send their entire position Q to DarkMid and execute it at zero cost. Therefore, even though they have executed the exact quantity Q , they have done so at a cost that is too high ex post. Conversely, conditional on only one investor being present, say the buyer, the buyer sends "too much" to DarkMid—"too much" because his DarkMid order is not executed at all. Put differently, if it turns out both investors are present, then they wish they had sent more to DarkMid; if it turns out only one investor is present, then the present investor wishes he had sent less to DarkMid.

news, or surprise earnings in the empirical sections. It raises investors' opportunity cost of not trading. One could argue that the elevated volatility following those shocks also raises inventory cost for intermediaries in DarkNMid and Lit, and therefore their opportunity cost of not trading is raised as well. The additional volume that follows these shocks, however, leads us to believe that at least part of the shock is attributable to a disproportionately large shock to investors as compared to intermediaries.

In calculating volume and market shares, we will focus on the $Q \geq \Delta$ case of Proposition 1, since in the data Lit venue has positive market share.

The expected volume in the three venues and the total volume are given by

$$\bar{v}_M = \phi^2(2x_M) + (1 - \phi)^2(2 \times 0) + 2\phi(1 - \phi)(2 \times 0) = \frac{2\phi^2\beta}{(1 - \phi)\gamma}, \quad (11)$$

$$\bar{v}_N = \phi^2(2x_N) + (1 - \phi)^2(2 \times 0) + 2\phi(1 - \phi)(x_N + 0) = \frac{2\phi\beta}{(1 - \phi)(\gamma + \eta)}, \quad (12)$$

$$\bar{v}_L = \phi^2(2x_L) + 2\phi(1 - \phi)x_L = 2\phi \left(Q - \frac{\beta}{(1 - \phi)\gamma} - \frac{\beta}{(1 - \phi)(\gamma + \eta)} \right), \quad (13)$$

$$\bar{v} = \bar{v}_M + \bar{v}_N + \bar{v}_L = 2\phi Q - \frac{2\phi\beta}{\gamma}. \quad (14)$$

Note that in the above calculation we double-count volume in DarkMid. In practice, operators of DarkMid typically act as buyer to the seller and seller to the buyer, so one match shows up as two trades. Our pecking order results are not affected if DarkMid volume is single-counted.

The volume shares of different venues are defined as,

$$s_i := \frac{\bar{v}_i}{\bar{v}}, \quad \text{for } i \in \{M, N, L\}. \quad (15)$$

Signing partial derivatives of volume shares with respect to γ yields the model's main proposition. It formalizes this paper's pecking order hypothesis, depicted in Panel (b) of Fig. 2.

Proposition 2 (Venue share and urgency). *As investor urgency increases, lit volume share increases and dark volume share decreases. Furthermore, DarkMid is more sensitive to urgency than DarkNMid:*

$$\frac{\partial s_M/s_M}{\partial \gamma/\gamma} < \frac{\partial s_N/s_N}{\partial \gamma/\gamma} < 0 < \frac{\partial s_L/s_L}{\partial \gamma/\gamma}. \quad (16)$$

This pecking order hypothesis is empirically supported by evidence shown in Section 5.

To understand what market conditions might affect the elasticities of venue market shares to urgency, we evaluate second-order partial derivatives of the venue shares s_i with respect to β and Q .

This exercise generates predictions on how strongly the pecking order hypothesis would show up in the cross-section of stocks. The result is stated in the following proposition.

Proposition 3 (Cross-section of venue share elasticities). *As the half-spread β in the lit exchange widens or as the trading interest Q decreases, the two dark venues' share elasticities (with respect to urgency γ) become more negative, while the lit elasticity becomes more positive. Mathematically,*

$$\frac{\partial}{\partial \beta} \left(\frac{\partial s_M}{\partial \gamma} \frac{\gamma}{s_M} \right) < 0, \quad \frac{\partial}{\partial \beta} \left(\frac{\partial s_N}{\partial \gamma} \frac{\gamma}{s_N} \right) < 0, \quad \text{and} \quad \frac{\partial}{\partial \beta} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) > 0; \quad (17)$$

$$\frac{\partial}{\partial Q} \left(\frac{\partial s_M}{\partial \gamma} \frac{\gamma}{s_M} \right) > 0, \quad \frac{\partial}{\partial Q} \left(\frac{\partial s_N}{\partial \gamma} \frac{\gamma}{s_N} \right) > 0, \quad \text{and} \quad \frac{\partial}{\partial Q} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) < 0. \quad (18)$$

Proposition 3 implies that the pecking order hypothesis should find stronger support in the data if spread is wider or if trading volume is smaller, which are characteristics of relatively illiquid stocks. This result can partly explain the lack of statistical significance for large stocks, as shown in Section 5.

8.2 A consolidated market benchmark

We have solved a model with a three-way volume fragmentation among Lit, DarkMid, and DarkNMid. The pecking order hypothesis follows from the model naturally. In this subsection we consider an alternative market structure with a single consolidated venue.

There are more than one way to model a consolidated market. Perhaps a trivial way is to consolidate DarkMid, DarkNMid and Lit into a single order book, in the following manner:

- Lit providers post unlimited depth at $\pm\beta$;
- The DarkNMid provider posts limit sell order at p_N and limit buy order at $-p_N$, as well as the associated quantities;
- The order book allows midpoint orders that can only be executed at the midpoint price, 0;
- The buyer chooses the size x_M of midpoint order, the size x_N of order that executes against the DarkNMid provider's limit orders, and the size x_L of market order to be executed at β .

Clearly, this consolidated market is equivalent to the fragmented market in the previous subsection.

A more interesting way to set up a consolidated market is to enforce a single trading mechanism. In particular, it would be interesting to let investors submit limit orders and make markets for each other. Toward this end, for the remainder of this subsection we will consider a consolidated market

that allows investors to post demand schedules, equivalent to a set of limit orders. This mechanism is similar to an open auction or a close auction on stock exchanges.

The asset and market participants are the same as in Section 8.1. The investors and the DarkNMid provider trade in a single Walrasian auction. If an investor is present, the investor submits a demand schedule $x(p)$. The demand schedule says that at the price p , the investor is willing to buy $x(p)$ units. A negative $x(p)$ is a sale. If an investor is not present, he of course cannot submit this demand schedule. The DarkNMid provider is always present and submits a demand schedule $y(p)$. At price $\pm\beta$, Lit liquidity providers are willing to accommodate any supply/demand imbalance, so the equilibrium price never goes outside of $[-\beta, \beta]$. Once the market-clearing price p^* is determined, the investors, the DarkNMid provider, and those Lit providers who execute some orders trade at the same price p^* . We refer to this consolidated market structure as “Lit++” (the double plus indicates the addition of both DarkMid and DarkNMid to Lit).

We assume that both investors and the DarkNMid provider take prices as given and maximize profit. Price-taking is a standard assumption in many models, such as rational expectations equilibrium (REE) models following Grossman and Stiglitz (1980).

We first solve the DarkNMid provider’s competitive price schedule. His problem is

$$\max_y -py - \frac{\eta}{2}y^2, \quad (19)$$

whose solution is

$$y(p) = -\frac{p}{\eta}, \quad p \in [-\beta, \beta]. \quad (20)$$

If the buyer is present, the buyer’s problem is

$$\max_x -px - \frac{\gamma}{2}(Q - x)^2, \quad (21)$$

which implies

$$x_B(p) = Q - \frac{p}{\gamma}, \quad p \in [-\beta, \beta]. \quad (22)$$

By the same argument, the seller’s demand schedule is $x_S(p) = -Q - p/\gamma$, $p \in [-\beta, \beta]$.

If both investors are present, the market clearing condition is $x_B(p) + x_S(p) + y(p) = 0$, which implies that the equilibrium price is 0 and both investors execute the full quantity Q .

If only one investor is present, the equilibrium price can be either interior (within the range $(-\beta, \beta)$) or corner (equal to β or $-\beta$), depending on Q . Without loss of generality, suppose that only the buyer is present. An interior p solves $0 = x_B(p) + y(p)$, or

$$p = \frac{Q}{1/\gamma + 1/\eta}. \quad (23)$$

The condition of an interior price $(-\beta) < p < \beta$ gives

$$Q < \beta \left(\frac{1}{\gamma} + \frac{1}{\eta} \right). \quad (24)$$

Conversely, a corner price of $p = \beta$ applies if

$$Q \geq \beta \left(\frac{1}{\gamma} + \frac{1}{\eta} \right). \quad (25)$$

At such a corner, $x_B = Q - \beta/\gamma$, the β/η of which is bought by the DarkNMid provider and the rest by the Lit providers. This completes the characterization of the equilibrium.

8.3 Comparing Lit++ to three-way fragmentation

In this subsection, we compare allocative efficiency of the two market structures. Since in reality all three venues have positive market shares, we will focus on sufficiently large Q so that the Lit providers execute some orders in both market structures. That is, we are working under the parameter assumption

$$Q \geq \frac{\beta}{1-\phi} \left(\frac{1}{\gamma} + \frac{1}{\gamma+\eta} \right) \text{ and } Q \geq \left(\frac{1}{\gamma} + \frac{1}{\eta} \right) \beta. \quad (26)$$

To measure allocative inefficiency, we need to compute the first-best allocation, which is achieved if every unit of the asset is distributed to whoever has the lowest marginal cost for holding it. There are two cases. When both investors are present, clearly, the first-best allocation is to let the buyer and the seller exchange their inventory entirely so that all three types of agents hold zero inventory. The aggregate inventory cost in this case is zero. When only one investor is present, for example the buyer, the first-best allocation under a large Q satisfying Eq. (26) is such that the marginal inventory costs of the buyer and the DarkNMid provider are equal to β . Thus, the buyer will hold β/γ units of the asset; the DarkNMid provider will hold β/η units; and the Lit providers hold the rest. The aggregate inventory cost in this case is

$$\frac{\gamma}{2}(\beta/\gamma)^2 + \frac{\eta}{2}(\beta/\eta)^2 + (Q - \beta/\gamma - \beta/\eta)\beta = \beta Q - \frac{\beta^2}{2} \left(\frac{1}{\gamma} + \frac{1}{\eta} \right). \quad (27)$$

Summing up the two cases, under Eq. (26), the expected aggregate inventory cost under the first-best

allocation is

$$k_0 = \phi^2 \cdot 0 + 2(1 - \phi)\phi \cdot \left[\beta Q - \frac{\beta^2}{2} \left(\frac{1}{\gamma} + \frac{1}{\eta} \right) \right] = 2(1 - \phi)\phi \cdot \left[\beta Q - \frac{\beta^2}{2} \left(\frac{1}{\gamma} + \frac{1}{\eta} \right) \right]. \quad (28)$$

In fact, we observe that the first-best allocation is achieved under the consolidated market structure. This is not surprising because the investors and the DarkNMid provider are price takers in the consolidated market; they adjust their holdings such that the market-clearing price is equal to the marginal inventory cost, guaranteeing full efficiency.

Under fragmentation, the total expected cost incurred by all the agents is

$$k_{MNL} = \phi^2 2x_L\beta + 2(1 - \phi)\phi \left(x_L\beta + \frac{\eta}{2}x_N^2 + \frac{\gamma}{2}x_M^2 \right) = 2\phi\beta Q - \frac{\phi}{1 - \phi}\beta^2 \left(\frac{1}{\gamma} + \frac{2\gamma + \eta}{(\gamma + \eta)^2} \right). \quad (29)$$

The allocative inefficiency of the fragmented market structure is then

$$k_{MNL} - k_0 = 2\phi^2\beta Q - \frac{\phi}{1 - \phi}\beta^2 \left(\frac{1}{\gamma} + \frac{2\gamma + \eta}{(\gamma + \eta)^2} \right) + (1 - \phi)\phi\beta^2 \left(\frac{1}{\gamma} + \frac{1}{\eta} \right), \quad (30)$$

which we can verify is positive under Eq. (26).

There is a simple intuition why allocations are less efficient under fragmentation. Since Lit providers have the highest inventory holding cost, it is desirable to use them only when necessary. The consolidated market has exactly this feature: Lit providers execute orders if and only if one investor is present, but not both. In contrast, under fragmentation, because an investor sends Lit orders without knowing if the other side is present, the Lit providers are used more often than necessary. Thus, the more efficient use of Lit providers favors consolidation for a sufficiently large Q . Likewise, the consolidated market structure uses the DarkNMid provider more efficiently (i.e., he executes orders if and only if one investor is present but not the other). This result captures the folk intuition that concentrating liquidity to a single venue enhances welfare.

We emphasize that the consolidated market achieves the first best because the investors and the DarkNMid providers take price as given, i.e., they do not strategically take into account their impact on the price. Dropping this price-taking assumption is likely to complicate the analysis substantially and lead to ambiguous welfare ranking. For example, Duffie and Zhu (2016) model the effect of adding a single “size discovery” mechanism, which matches buy and sell orders at a fixed price, to a sequential double auction market with strategic investors. Because strategic investors mitigate their own price impact by trading less aggressively than competitive investors, their consolidated market (sequential double auctions) does not achieve the first best. However, adding a size discovery mechanism, which freezes the price, improves allocative efficiency because investors trade more aggressively without concern of price impact. If we were to repeat their analysis in our context, we would expect the same

benefit of fragmentation due to DarkMid, hence an ambiguous welfare ranking between consolidation and fragmentation.

9 Conclusion

We propose and test a pecking order hypothesis for the dynamic fragmentation of U.S. equity markets. The hypothesis posits that investors disperse their orders across various venue types according to a pecking order. The position of venue types on the pecking order depends on the tradeoff between trading cost (price impact) and immediacy (execution certainty). On top of the pecking order are the low-cost, low-immediacy venues such as midpoint dark pools (DarkMid), whereas at the bottom are high-cost, high-immediacy venues such as lit exchanges (Lit); in the middle of the pecking order are non-midpoint dark pools (DarkNMid). A positive shock to investors' urgency to trade tilts their order flow from the top of the pecking to the bottom; therefore, the elasticities of venue market shares to urgency shocks are progressively less negative and more positive further down the pecking order. We micro-found the pecking order hypothesis with a stylized model of strategic venue choice among DarkMid, DarkNMid, and Lit.

We test the pecking order hypothesis using a unique U.S. equity dataset that, for each trade, identifies in what type of venue it transpired. The dataset distinguishes five venue types: two types of dark pools, namely DarkMid and DarkNMid, and three types of broker-dealer internalization, namely retail trades, average-price trades, and other trades. We use three exogenous variables as proxies for investors' urgency to trade: VIX, macroeconomic data releases, and earnings announcements. These three variables are used in a panel VARX model that characterizes the dynamic interrelation between dark trading volumes and (endogenous) measures of market conditions such as spread, depth, and volatility.

As predicted by the pecking order hypothesis, a positive shock to VIX substantially reduces the market share of DarkMid, moderately reduces the market share of DarkNMid, but increases the market share of Lit. Macroeconomic data releases show a similar but stronger pattern. After surprise earnings announcements, the share of midpoint pools also declines significantly. After disaggregating across by size, we find strong support for the pecking order hypothesis for small and medium stocks, but no significant pattern for large stocks. This weak result for large stocks is predicted by our stylized model.

Appendix

A Snippet of disaggregated dark transaction data

The following table provides a snippet of the raw transaction data of Alcoa, disaggregated to five types of dark venues. The bolded field has the five categories explained in the text: DarkMid (MP), DarkNMid (DP), DarkRetail (RT), DarkPrintB (PB), and DarkOther (OT).

date	time	symbol	type	contra	buysell	price	shares	cond1	cond2	cond3	cond4
1-Oct-10	9301833	AA	DP	BD	B	12.2875	100	@			
1-Oct-10	9301834	AA	DP	BD	B	12.2875	100	@			
1-Oct-10	9301941	AA	OT	BD	S	12.28	100	@			
1-Oct-10	9301989	AA	DP	BD	B	12.285	200	@			
1-Oct-10	9302005	AA	MP	BD	X	12.285	100	@			
1-Oct-10	9302148	AA	RT	BD	S	12.29	9000	@			
1-Oct-10	9302204	AA	RT	BD	B	12.2701	300	@			
1-Oct-10	9302224	AA	RT	BD	B	12.27	100	@			
1-Oct-10	9302249	AA	DP	BD	B	12.28	100	@			
1-Oct-10	9302343	AA	DP	BD	B	12.28	100	@			
1-Oct-10	9302540	AA	DP	BD	B	12.2805	100	@			
1-Oct-10	9302546	AA	RT	BD	S	12.29	160	@			
...
1-Oct-10	10100150	AA	PB		S	12.285	179379	@	4		B

B Details on the implementation of the panel VARX model

In this appendix we discuss the details of the panel VARX model. The estimation is implemented via OLS by stacking the observations associated with different stocks into a single vector. The stock fixed effect is accounted for by adding dummy variables to the set of regressors. Lags of the variables are only constructed intraday. Time-of-day dummies (13 in total, one for each half-hour trading interval) are also included for each stock.

The optimal number of lags p is chosen according to Bayesian Information Criterion (BIC). Specifically, for each of the 117 stocks, the VARX model is estimated for all $p \in \{1, 2, \dots, 10\}$. Then the best (according to BIC) lag is chosen for stock j . That is, we confine the search of the optimal lags within 10 lags for the endogenous variables. The above procedure generates 117 optimal lags of $\{p_j\}_{j=1}^{117}$. There are 16 p_j that are found to be 1, 93 to be 2, and the other 8 to be 3. We hence choose $p = 2$ for parsimony.

The standard errors for panel data estimators should account for potential correlation through time and across stocks. One standard way to account for these issues is to do “double-clustering” (Petersen, 2009). The laborious (but most flexible) way of implementing such clustering is by calculating

$$\text{cov}(\hat{\beta}_i, \hat{\beta}_j) = (X'X)^{-1} V (X'X)^{-1} \text{ with } v_{ij} = \sum_{kt,ls} x_{ikt} \hat{\epsilon}_{ikt} \hat{\epsilon}_{jls} x_{jls} \times 1_A(ktls), \quad (31)$$

where $i, j \in \{1, \dots, N\}$ where N is the number of regressors, $k, l \in \{1, \dots, J\}$ where J is the number of stocks, and $s, t \in \{1, \dots, T\}$ where T is the number of time periods. $1_A(ktls)$ is the indicator function where the subset A of the index value space identifies which auto- or cross-correlations a researcher worries about. If error terms are independent and identically distributed, then the indicator function equals one if $k = l$ and $t = s$. The subset A for an indicator function in a standard double-clustering is such that:

$$1_A(ktls) = \begin{cases} 1 & \text{if } k = l \text{ or } s = t, \\ 0 & \text{otherwise.} \end{cases} \quad (32)$$

A researcher can easily be more conservative and also account for non-zero cross-autocorrelations by also including changing the $s = t$ condition by, say, $|s - t| \leq 5$.

The cumulative impulse response function is most easily calculated by stacking the estimated Φ matrices, as any VAR can always be written as a first-order VAR. Consider, for example, a VAR with two lags. This VAR can be written as

$$\begin{bmatrix} y_t \\ y_{t-1} \end{bmatrix} = \begin{bmatrix} \Phi_1 & \Phi_2 \\ I & 0 \end{bmatrix}' \begin{bmatrix} y_{t-1} \\ y_{t-2} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix}. \quad (33)$$

The t -period cumulative impulse response of the j^{th} variable to a unit impulse in the i^{th} variable is the j^{th} element of the vector

$$\begin{bmatrix} \Phi_1 & \Phi_2 \\ I & 0 \end{bmatrix}' e_k, \quad (34)$$

where I is the identity matrix and e_k is the unit vector where the k^{th} element is one and all other elements are zeros. When an exogenous variable is shocked, one can simply scale the shock size by the contemporaneous responses of endogenous variables to the exogenous shock by resorting to the estimates Ψ .

Confidence intervals on the impulse response function (IRF) are obtained through simulation. The IRF is a non-linear transformation of the VARX coefficient estimates denoted by $\hat{\theta}$. Each simulation involves a draw from the multivariate normal distribution $N(\hat{\theta}, \Sigma_{\hat{\theta}})$, where $\Sigma_{\hat{\theta}}$ is the estimated double-clustered covariance matrix of the coefficients. Note that this distribution is asymptotically true given the assumption that the VARX model is correctly specified with normal residual terms. We perform 10,000 independent draws of the coefficients $\hat{\theta}$ and for each draw compute the IRFs at all lags. Thus, we obtain, for each IRF, an i.i.d. sample of size 10,000. The confidence bounds are then chosen at the 2.5 and the 97.5 (or 0.5 and 99.5) percentiles of the simulated IRFs. The significance levels shown in Table 4 are based on whether the estimates exceed the confidence bounds found above.

C Transformation between logarithms and levels of market share variables

In the VARX model implementation, the variables are log-transformed (except *EpsSurprise*). Log-transformation has several advantages. For example, the strictly positive variables (e.g. volume,

spread, depth, etc.) are converted to a possibly negative support; the concavity in logarithm discourages the abnormal effects of outliers; the estimation coefficients can be readily interpreted as elasticity. The key variables of our focus are the (logged) trading volumes in the five dark venues and the lit venue, denoted by $\log v_1, \dots, \log v_5$, and $\log v_6$, where the first five are for the five dark venues and the last v_6 is for the trading volume in the lit. (Each of these variables has stock-day-minute granularity.) For the pecking order hypothesis, it is however useful to think in terms of market shares, defined as

$$s_j = \frac{v_j}{\sum_j v_j}$$

for $j \in \{1, \dots, 6\}$.

The purpose of this appendix is to derive the closed-form, exact transformation formula from a shock in trading volume in one venue to the response of all market shares. Specifically, given a shock of $\Delta \log v_i$, we want to know the immediate response Δs_j , for *all* $j \in \{1, \dots, 6\}$. Reverse directions from Δs_i to $\Delta \log v_j$ will also be dealt with. These formulas are used in generating the impulse responses in testing the pecking order hypothesis.

In the derivation below, we shall use the following additional notations. Let v be the total volume: $v = \sum_j v_j$. We shall use a superscript of “+” to denote the variables after a shock; for example, $\log v_j^+ = \log v_j + \Delta \log v_j$. Similarly, while s_j denotes the market share of venue j , s_j^+ denotes the market share after the shock. In the IRF exercise, the pre-shock values will be chosen as the stock-day-minute average across all raw sample observations. Consider the following cases.

From $\Delta \log v_i$ to Δs_j . By construction, $\log v_i^+ = \log v_i + \Delta \log v_i$. Taking the exponential on both sides gives the *level* of the post-shock trading volume: $v_i^+ = v_i \exp \Delta \log v_i$. The post-shock market share by construction is

$$s_i^+ = \frac{v_i^+}{v^+} = \frac{v_i^+}{v + (v_i^+ - v_i)}.$$

Substituting with the expression of v_i^+ and then subtracting $s_i = v_i/v$ yields

$$\Delta s_i = s_i^+ - s_i = \dots = s_i \cdot \left(e^{\Delta \log v_i - \Delta \log v} - 1 \right) \quad (35)$$

where $\Delta \log v = \log v^+ - \log v = \log(\sum_{j \neq i} v_j + v_i^+) - \log v$. The above formula actually applies to both the venue i whose volume is shocked and any other venue $j \neq i$ whose volume is not shocked. The only difference is, as can be seen after substituting the index i with a different j , that $\Delta \log v_j = 0$ for $j \neq i$. Finally, we can immediately derive the dark volume share change as the complement of the change in the lit share: $\sum_{j \leq 5} \Delta s_j = -\Delta s_6$, simply because the identity of $s_6 = 1 - \sum_{j \leq 5} s_j$.

From $\Delta \log v$ to Δs_i , assuming proportionally scaling across all venues. Now we shock the total volume such that $\log v^+ = \log v + \Delta \log v$ and make the assumption that the increase in volume is proportionally scaled across all venues. That is, for each venue i , $v_i^+ - v_i = s_i \cdot (v^+ - v)$, or

$$v_i^+ = v_i + s_i \Delta v = s_i v + s_i \Delta v = s_i \cdot (v + \Delta v) = s_i v^+.$$

Taking logarithm on both sides gives $\log v_i^+ = \log s_i + \log v^+$. Substitute $\log s_i$ with $\log s_i = \log(v_i/v) = \log v_i - \log v$ and then

$$\log v_i^+ = \log v_i - \log v + \log v^+ \implies \Delta \log v_i = \Delta \log v. \quad (36)$$

Applying Eq. (36) to Eq. (35) immediately gives $\Delta s_i = 0$. Clearly, this holds for all $i \in \{1, \dots, 6\}$.

From Δs_i to $\Delta \log v_j$ by shocking $\log v_i$ and proportionally offsetting in other venues, without changing total volume v . Finally we do the reverse. Suppose we shock s_i by Δs_i . Such a change in market share must be driven by some change(s) in trading volume(s). Here a particular change is considered: Let v_i change in the same direction as s_i but all other $v_{j \neq i}$ move in the other direction so that the total volume does not change, i.e. $v = v^+$. We want to know, given the size of Δs_i , what are the sizes of $\Delta \log v_j$ for all $j \in \{1, \dots, 6\}$.

First, consider $j = i$. By construction, $\Delta s_i = v_i^+/v^+ - s_i$. Because the total volume is assumed to be unchanged, we have $v_i^+ = (s_i + \Delta s_i)v$. This enables the second equality below:

$$\begin{aligned} \Delta \log v_i &= \log v_i^+ - \log v_i \\ &= \log(s_i + \Delta s_i) + \log v - \log v_i = \log(s_i + \Delta s_i) + \log \frac{v}{v_i} = \log(s_i + \Delta s_i) - \log s_i = \log \left(1 + \frac{\Delta s_i}{s_i} \right). \end{aligned} \quad (37)$$

Consider next $j \neq i$. To offset Δv_i , summing over all $j \neq i$ gives $\sum_{j \neq i} \Delta v_j = -\Delta v_i$. Because the changes are proportional according to s_j , we have

$$v_j^+ = v_j - \frac{s_j}{\sum_{h \neq i} s_h} \Delta v_i = s_j \cdot \left(v - \frac{\Delta v_i}{\sum_{h \neq i} s_h} \right).$$

Take logarithm on both sides and expand $\log s_j = \log v_j - \log v$ to get

$$\log v_j^+ - \log v_j = \Delta \log v_j = -\log v + \log \left(v - \frac{\Delta v_i}{\sum_{h \neq i} s_h} \right) = \log \left(1 - \frac{\Delta v_i/v}{\sum_{h \neq i} s_h} \right) = \log \left(1 - \frac{\Delta s_i}{\sum_{h \neq i} s_h} \right), \quad (38)$$

where the last equality follows because the total volume is assumed to be unchanged: $\Delta v_i/v = v_i^+/v - v_i/v = v_i^+/v^+ - v_i/v = s_i^+ - s_i = \Delta s_i$.

D Proofs

D.1 Proof of Proposition 1

By symmetry, we only focus on the buyer's strategy.

Interior solution. To begin with, suppose the equilibrium DarkNMid price is interior, i.e. $p_N \in [0, \beta)$. This means that the buyer prefers not to trade in Lit, where the marginal cost of trading is β .

Splitting Q across DarkMid and DarkNMid means that in equilibrium, the marginal costs equate each other:

$$(1 - \phi)\gamma x_M = p_N. \quad (39)$$

Using $Q = x_M + x_N$, we have

$$x_N = Q - \frac{p_N}{(1 - \phi)\gamma}. \quad (40)$$

We conjecture, and later verify, that the DarkNMid provider is willing to execute the entire quantity x_N . We then solve the equilibrium p_N set by the DarkNMid provider. Recall that the DarkNMid liquidity provider maximizes profit, which is

$$\Pi = 2\phi p_N x_N - 2\phi(1 - \phi)\frac{1}{2}\eta x_N^2 = 2\phi p_N \left(Q - \frac{p_N}{(1 - \phi)\gamma} \right) - \phi(1 - \phi)\eta \left(Q - \frac{p_N}{(1 - \phi)\gamma} \right)^2. \quad (41)$$

Taking the first-order condition with respect to p_N , we can solve

$$p_N = \frac{\gamma(\gamma + \eta)}{2\gamma + \eta}(1 - \phi)Q. \quad (42)$$

An interior solution requires $p_N < \beta$, or

$$Q < \Delta := \frac{\beta}{1 - \phi} \left(\frac{1}{\gamma} + \frac{1}{\gamma + \eta} \right). \quad (43)$$

Under the above interior p_N , we immediately have

$$x_M = \frac{\gamma + \eta}{2\gamma + \eta}Q, \quad x_N = \frac{\gamma}{2\gamma + \eta}Q, \quad x_L = 0. \quad (44)$$

We now verify that, at p_N , the DarkNMid provider does not wish to execute less than x_N . Fixing the quotes to the seller, consider a deviation of the quote to the buyer to (p_N, q) , with $q \leq x_N$. The DarkNMid provider's profit made from the buyer is then

$$\phi p_N q - \phi(1 - \phi)\frac{1}{2}\eta q^2, \quad (45)$$

which implies that the marginal profit of increasing q is $\phi(p_N - (1 - \phi)\eta q)$. Substitute in the solution p_N and $q = x_N$, we see that the marginal profit of increasing quantity is ϕ times:

$$\frac{\gamma(\gamma + \eta)}{2\gamma + \eta}(1 - \phi)Q - (1 - \phi)\eta \frac{\gamma}{2\gamma + \eta}Q > 0. \quad (46)$$

That is, the DarkNMid provider actually wishes to execute more than x_N at p_N . This completes the characterization of the interior solution of the equilibrium.

Corner solution. At the corner solution, $p_N = \beta$. Equating the marginal cost in DarkMid, $(1 - \phi)\gamma x_M$, to β gives $x_M = \beta/((1 - \phi)\gamma)$.

To solve x_L and x_N , consider the decision of the buyer to split between DarkMid and DarkNMid after sending x_L to Lit. Using the same argument as above, the optimal x_N chosen by the buyer at the price $p_N \leq \beta$ is

$$x_N = Q - x_L - \frac{p_N}{(1-\phi)\gamma}. \quad (47)$$

Again, conjecture (and later verify) that the DarkNMid provider is willing to execute the full quantity x_N . Substituting the above x_N into profit Π and taking the first-order condition, we have

$$p_N = \frac{\gamma(\gamma + \eta)}{2\gamma + \eta}(1-\phi)(Q - x_L). \quad (48)$$

Since $p_N = \beta$ is the optimal solution, we have

$$Q - x_L = \frac{\beta}{1-\phi} \frac{2\gamma + \eta}{\gamma(\gamma + \eta)} = x_M + x_N = \frac{\beta}{(1-\phi)\gamma} + x_N, \quad (49)$$

from which we get

$$x_N = \frac{\beta}{(1-\phi)(\gamma + \eta)} \quad (50)$$

Obviously, $x_L \geq 0$ implies $Q \geq \Delta$. Moreover, as in the interior solution, we verify that, at $p_N = \beta$, the marginal profit of increasing quantity is ϕ times:

$$\beta - (1-\phi)\eta \cdot \frac{\beta}{(1-\phi)(\gamma + \eta)} > 0. \quad (51)$$

So the DarkNMid provider executes the full quantity x_N at β .

D.2 Proof of Proposition 2

Direct calculation shows:

$$s_M = \frac{\phi\beta}{(\gamma Q - \beta)(1-\phi)}, \quad (52)$$

$$s_N = \frac{\beta\gamma}{(\gamma Q - \beta)(\gamma + \eta)(1-\phi)}, \quad (53)$$

$$\frac{\partial s_M}{\partial \gamma} = -\frac{\phi\beta Q}{(\gamma Q - \beta)^2(1-\phi)} < 0, \quad (54)$$

$$\frac{\partial s_N}{\partial \gamma} = -\frac{\beta(\gamma^2 Q + \beta\eta)}{(\gamma Q - \beta)^2(\gamma + \eta)^2(1-\phi)} < 0, \quad (55)$$

$$\frac{\partial s_L}{\partial \gamma} = 1 - \frac{\partial s_M}{\partial \gamma} - \frac{\partial s_N}{\partial \gamma} > 0, \quad (56)$$

where the last line follows from $s_M + s_N + s_L = 1$.

It remains to rank the volume share sensitivity to urgency of the two dark venues. Direct calculation shows:

$$\frac{\partial s_M}{\partial \gamma} \frac{\gamma}{s_M} - \frac{\partial s_N}{\partial \gamma} \frac{\gamma}{s_N} = -\frac{\eta}{\gamma + \eta} < 0. \quad (57)$$

Hence, the sensitivities of volume shares to urgency can be ranked as stated in Proposition 2.

D.3 Proof of Proposition 3

It is easy to compute the venue share elasticities for the two dark venues:

$$\frac{\partial s_M}{\partial \gamma} \frac{\gamma}{s_M} = -\frac{\gamma Q}{\gamma Q - \beta} < 0 \quad (58)$$

$$\frac{\partial s_N}{\partial \gamma} \frac{\gamma}{s_N} = -\frac{\beta}{\gamma Q - \beta} - \frac{\gamma}{\gamma + \eta} < 0. \quad (59)$$

Hence, these two elasticities are increasing in Q and decreasing in β .

The venue share elasticity for the lit venue has the following expression:

$$\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} = \frac{\beta(\beta\eta + Q(\gamma^2 + (\gamma + \eta)^2\phi))}{(\gamma Q - \beta)(\gamma + \eta)^2(1 - \phi)(Q - \Delta)}. \quad (60)$$

The cross-derivatives are:

$$\frac{\partial}{\partial \beta} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) = \gamma Q \frac{2Q\beta\gamma^2\eta(1 - \phi) + Q^2\gamma^2(1 - \phi)(\gamma^2 + (\gamma + \eta)^2\phi) - \beta^2(\eta^2\phi + 2\gamma^2(1 + \phi) + 2\gamma\eta(1 + \phi))}{(\gamma Q - \beta)^2(Q\gamma(\gamma + \eta)(1 - \phi) - \beta(2\gamma + \eta))^2}, \quad (61)$$

$$\frac{\partial}{\partial Q} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) = -\beta\gamma \frac{2Q\beta\gamma^2\eta(1 - \phi) + Q^2\gamma^2(1 - \phi)(\gamma^2 + (\gamma + \eta)^2\phi) - \beta^2(\eta^2\phi + 2\gamma^2(1 + \phi) + 2\gamma\eta(1 + \phi))}{(\gamma Q - \beta)^2(Q\gamma(\gamma + \eta)(1 - \phi) - \beta(2\gamma + \eta))^2}. \quad (62)$$

So it only remains to sign the common numerator of these two cross derivatives. Note that this numerator is quadratic increasing in Q . Its minimum is achieved if $Q = \Delta$. Substituting Δ into this numerator yields a minimum of

$$\frac{\beta^2(2\gamma^2 + 2\gamma\eta + \eta^2)(\gamma + (\gamma + \eta)\phi)^2}{(\gamma + \eta)^2(1 - \phi)} > 0, \quad (63)$$

implying that the numerator is strictly positive on the support of $Q \geq \Delta$. Therefore,

$$\frac{\partial}{\partial \beta} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) > 0 \text{ and } \frac{\partial}{\partial Q} \left(\frac{\partial s_L}{\partial \gamma} \frac{\gamma}{s_L} \right) < 0. \quad (64)$$

This completes the proof.

References

- Barber, B. M., Odean, T., 2007. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* pp. 785–818.
- Bloomfield, R., O'Hara, M., Saar, G., 2015. Hidden liquidity: Some new light on dark trading. *Journal of Finance* 70, 2227–2273.
- Boni, L., Brown, D. C., Leach, J. C., 2012. Dark pool exclusivity matters. Working paper.
- Boulatov, A., George, T. J., 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. *Review of Financial Studies* 26, 2095–2137.
- Brogaard, J., Hendershott, T., Riordan, R., 2014. High frequency trading and price discovery. *Review of Financial Studies* 27, 2267–2306.
- Brolley, M., 2014. Should dark pools improve upon visible quotes? the impact of trade-at rules. Working paper, University of Toronto.
- Buti, S., Rindi, B., Werner, I. M., 2011. Diving into dark pools. Working paper, Fisher College of Business, Ohio State University.
- Buti, S., Rindi, B., Werner, I. M., 2015. Dark pool trading strategies, market quality and welfare. *Journal of Financial Economics* (Forthcoming) .
- Campbell, J. Y., Grossman, S. J., Wang, J., 1993. Trading volume and serial correlation in stocks returns. *The Quarterly Journal of Economics* 108, 905–939.
- Chowdhry, B., Nanda, V., 1991. Multimarket trading and market liquidity. *Review of Financial Studies* 4, 483–511.
- Comerton-Forde, C., Malinova, K., Park, A., 2015. The impact of the dark trading rules. Report prepared for the investment industry regulatory organization of Canada.
- Comerton-Forde, C., Putniņš, T., 2015. Dark trading and price discovery. *Journal of Financial Economics* 118, 70–92.
- Degryse, H., de Jong, F., van Kervel, V., 2015a. The impact of dark trading and visible fragmentation on market quality. *Review of Finance* 19, 1587–1622.
- Degryse, H., Tombeur, G., Wuyts, G., 2015b. Two shades of opacity: Hidden orders versus dark trading. Working paper.
- Degryse, H., Van Achter, M., Wuyts, G., 2009. Dynamic order submission strategies with competition between a dealer market and a crossing network. *Journal of Financial Economics* 91, 319–338.
- Devani, B., Anderson, L., Zhang, Y., 2015. Impact of the dark rule amendments. Trading review and analysis, Investment Industry Regulatory Organization of Canada.

- Duffie, D., Zhu, H., 2016. Size discovery. Working paper, NBER.
- Ellis, K., Michaely, R., O'Hara, M., 2000. The accuracy of trade classification rules: Evidence from nasdaq. *Journal of Financial and Quantitative Analysis* 35, 529–551.
- Foley, S., Malinova, K., Park, A., 2013. Dark trading on public exchanges. Working paper.
- Foley, S., Putniņš, T., 2015. Should we be afraid of the dark? dark trading and market quality. Working paper.
- Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70(3), 393–408.
- Hatheway, F., Kwan, A., Zheng, H., 2013. An empirical analysis of market segmentation on u.s. equities markets. Working paper.
- Hendershott, T., Jones, C. M., 2005. Island goes dark: Transparency, fragmentation, and regulation. *Review of Financial Studies* 18, 743–793.
- Hendershott, T., Mendelson, H., 2000. Crossing networks and dealer markets: Competition and performance. *Journal of Finance* 55, 2071–2115.
- Holden, C. W., Jacobsen, S., 2014. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance* 69, 1747–1785.
- Kim, O., Verrecchia, R. E., 1994. Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics* 17, 41–67.
- Kinney, W., Burgstahler, D., Martin, R., 2002. Earnings surprise “materiality” as measure by stock returns. *Journal of Accounting Research* 40, 1297–1329.
- Kwan, A., Masulis, R., McInish, T., 2015. Trading rules, competition for order flow and market fragmentation. *Journal of Financial Economics* 115, 330–348.
- Kyle, A. S., 1989. Informed speculation with imperfect competition. *The Review of Economic Studies* 56, 317–355.
- Lee, C., Ready, M., 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733–746.
- Nimalendran, M., Ray, S., 2014. Informational linkages between dark and lit trading venues. *Journal of Financial Markets* 17, 230–261.
- O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? *Journal of Financial Economics* 100, 459–474.
- Pagano, M., 1989. Trading volume and asset liquidity. *Quarterly Journal of Economics* 104, 255–274.
- Parlour, C. A., 1998. Price dynamics in limit order markets. *The Review of Financial Studies* 11, 789–816.

- Parlour, C. A., Seppi, D. J., 2008. Limit order markets: A survey. In: Boot, A., Thakor, A. (eds.), *Handbook of Financial Intermediation and Banking*, Elsevier Publishing, Amsterdam, Netherlands.
- Petersen, M. A., 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* 22, 435–480.
- Ready, M. J., 2014. Determinants of volume in dark pool crossing networks. Working paper, University of Wisconsin-Madison.
- Sarkar, A., Schwartz, R. A., 2009. Market sidedness: Insights into motives for trade initiation. *Journal of Finance* 64, 375–423.
- SEC, 2010. Concept release on equity market structure. Release No. 34-61358; File No. S7-02-10.
- Seppi, D. J., 1997. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies* 10, 103–150.
- Tuttle, L., 2014. Otc trading: Description of non-ats otc trading in national market system stocks. Memorandum, SEC.
- Ye, M., 2011. A glimpse into the dark: Price formation, transaction costs and market share of the crossing network. Working paper, University of Illinois.
- Zhu, H., 2014. Do dark pools harm price discovery? *Review of Financial Studies* 27, 747–789.