Equilibrium Bid-Price Dispersion[1]

Boyan Jovanovic and Albert J. Menkveld

August 5, 2020

[1]Boyan Jovanovic, Department of Economics, New York University, 19 West 4th Street, New York, United States, tel +1 212 998 8953, bj2@nyu.edu. Albert J. Menkveld, Vrije Universiteit Amsterdam, SBE, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands, tel +31 20 598 6130, albertjmenkveld@gmail.com and Tinbergen Institute. We thank Kenneth Burdett, Benjamin Connault, Darrell Duffie, and conference participants at the tenth Philadelphia Search and Matching conference for valuable comments. We thank the NSF and NWO for support and Maikol Cerda, Sean Flynn, Seher Gupta, Björn Hagströmer, Wenqian Huang, Sai Ma, Gaston Navarro, Hargungeet Singh, Walter Verwer, Dong Wei, Xi Xiong, and Shihao Yu for assistance. Albert is grateful to NWO for a Vici grant.

Equilibrium Bid-Price Dispersion

Abstract

If bidding in a common-value auction is costly and if bidders do not know how many others are also bidding, all equilibria are in mixed strategies. Participation is probabilistic and bid prices are dispersed. The symmetric equilibrium is unique and yields simple analytic expressions. We use them to, for example, show that bid prices exhibit negative skewness. The expressions are further used to estimate the model based on bidding on an S&P500 security. We find that the number of bidders declined over time, making liquidity supply fragile.

Please find online appendix here: https://bit.ly/2BlV3Lf.

1 Introduction

There has been widespread concern about market liquidity in the aftermath of the 2007–2009 financial crisis. New regulation such as BIS III and Dodd-Frank aim to reduce risk in the banking sector. The Volcker Rule, part of Dodd-Frank, specifically affects market liquidity that severely constrains market-making by banks. Only if there is a demonstrable interest of their clients can banks participate as middlemen in the market.[1] BIS (2014) report a decline in dealer risk-taking capacity and/or willingness and diminished proprietary trading by banks (which could be considered market making by third parties (Duffie, 2012)).

Empirical studies show that the nature of market liquidity in the years since the crisis seems to have changed. The bid-ask spread remained relatively flat, yet large orders seem to have become more costly to execute and markets are more prone to "flash crashes." Standard costly-inventory (or adverse-selection) models might explain the increased cost, but they do not permit flash crashes (e.g., Grossman and Miller, 1988).[2]

We propose a costly bidding model that could generate these two empirical findings through endogenous middleman participation. In the model, homogeneous players decide simultaneously on bidding in a first-price auction. If they incur an (opportunity) cost to place a bid, then there is no equilibrium in pure strategies. There is, however, a unique symmetric equilibrium in mixed strategies. In this equilibrium, each middleman tosses a biased coin to decide whether or not to bid. If he bids, then he pays the cost and draws a price from a particular distribution. We would like to emphasize that simultaneous

[1] Duffie (2012) provides an economic perspective on how the Volcker Rule affects market making. Federal Register (2014) provides the details of the final rule and a summary of feedback by market participants.

[2] Adrian et al. (2017) survey the literature on bond-market liquidity after the financial crisis. Most studies find that the bid-ask spread mostly recovered soon after the crisis (e.g., Trebbi and Xiao, 2019; Anderson and Stulz, 2017; Bessembinder et al., 2018). Anderson and Stulz (2017) however, note that for large (>$100,000) orders transaction cost remain elevated. Dick-Nielsen (2019) and Bao, O'Hara, and Zhou (2018) find the same for stressed sells. Choi and Huh (2017) and Bessembinder et al. (2018) report that dealers participate less as principal. Finally, aside from concern about the liquidity level, there is concern about liquidity *risk* in the sense of increased intensity of illiquidiy spikes (see, e.g., Adrian et al. (2015) for evidence and Kennedy et al. (2015) for a review of the public debate). A recent article argues that even key futures markets show signs of fragility: "Thinning Liquidity in Key Futures Market Worries Traders," Financial Times, March 25.

bidding implies that bidders do not observe potential bids of others.

The equilibrium yields surprisingly simple analytic expressions. Both the probability of bidding and the bid-price distribution are obtained as a function of the three model parameters: the value of winning the auction (v), the cost of bidding (c), and the reservation value of the seller (u).

The equilibrium bid-price distribution is like an iceberg, the tip is always there but what is underneath the water remains uncertain. Let us clarify by example. Suppose that the bid-price support (i.e., $v - c - u$) shrinks but the bidding cost (i.e., c) remains unchanged, then the equilibrium bid-price distribution changes as follows. The *best* bid, the tip of the iceberg, will not change much as all its quantiles remain unchanged except for the very lowest ones (see Figure 1, which plots the equilibrium best bid distribution when the bid-price support shrinks as u increases from u_2 to u_1). Yet, on average, fewer middlemen participate resulting in fewer bids and therefore less depth, less ice underneath the water. And, importantly, the fragility of liquidity supply increases in the sense of an increased probability of no bids at all. The change in this example is more or less what we find in the data. Indeed, the level of the best bid is rather stable, yet the supply of bids has become more erratic over time thus consistent with the empirical evidence on post-crisis liquidity (see summary in footnote 2).

Another unequivocal prediction of the model is a negative skewness in bid prices. More precisely, the equilibrium bid-price density is increasing and convex and therefore exhibits negative skewness. The intuition is as follows. When a player (who decided to bid) considers what price to bid at, he understands that a lower price will yield a higher profit when winning but, at the same time, it is less likely to win as it becomes more likely that others entered with a higher bid. The probability of winning, however, drops disproportionately for ever lower prices, thus inducing stronger price shading further out in the left tail (to keep the player indifferent between candidate bid prices).

The model is taken to the data to test one of its main predictions, convexity in the bid-price density, and to estimate its parameters. The sample consists of bid-price snapshots for SPY from 2007 through 2018. SPY is an exchange-traded fund (ETF) that tracks the S&P500 and is one of the most actively traded securities.

The model is estimated by matching three empirical moments that uniquely determine the three parameters. These moments are:

1. *Bid aggressiveness.* Bid aggressiveness is measured as the ratio of bid-price densities at the top and at the bottom of a pre-specified price interval. In the model this ratio is larger than one because the bid-price density is upward sloping.

2. *Variation in the number of bids.* Time variation in the number of bids is measured by the *coefficient of variation* (CV) which is defined as the standard deviation divided by the mean. The use of this standard statistic is novel in the empirical literature on liquidity supply. We believe it is an important one to include when considering fragility.

3. *Bid-ask spread.* The relative bid-ask spread is the distance between the best bid and the best ask price, scaled by the midquote (i.e., the middle of the best bid and the best ask).

These three moments identify the model parameters (v, c, u) in the following way. The first moment, bid aggressiveness, uniquely determines the value of winning v via the shape of the distribution function. The second and third moment jointly identify the two remaining parameters: c and u. The second moment, the variation in the number of bids, is inversely related to the expected number of middlemen showing up to post a bid. This relationship is intuitive as high variation through time coincides with fewer middlemen showing up on average. In the model this average is a function of all three parameters and, therefore, one more moment is needed to identify the two remaining parameters c and u separately.

The third moment, the bid-ask spread, delivers full identification. The intuition is that the spread is tightly related to c because its lower bound is $2c$ (assuming symmetry across the bid and ask side of the book). The reason is that middlemen must earn at least c to make up for the cost of bidding. The identification of u then mostly comes from the second moment because, given a level of c, a lower u implies a larger bid-price support $[u, v-c]$ and therefore more opportunity for middlemen to participate, which, in equilibrium, results in more participation (i.e., a higher second moment). In summary, the three moments together identify the three parameters uniquely.

And, for any given level of c, a lower u implies a larger bid-price support $[u, v-c]$ and therefore more ground for middlemen to participate leading to higher expected participation in equilibrium. Therefore c and u are identified based on the spread (i.e., third moment) along with the average number of middlemen (i.e., second moment). In summary, the three moments together identify the three parameters uniquely.

Results. The empirical analysis yields the following insights. First, the predicted convexity of bid-price densities is a robust feature of bid-price dispersions in the data.

Second, the three moments used to estimate the model parameters show the following trends. Bid aggressiveness, the first moment, steadily increased with peaks around 2013-2014 (the period during which Dodd-Frank regulation was put in force which made market-making more costly for banks). Viewed through the lens of the model, such a trend implies that the value of winning v decreased over time with the lowest values attained in these two years.

The second moment, variation in the supply of bids, steadily increased over time growing 15.5% annually. The model can generate such trend by a decline in the expected participation of middlemen. The estimates suggest that around 40 middlemen participated at the start of the sample with less than half of this number towards the end of it. Note that this is a statement about *average* participation. It could mean the same number of market makers, *ex-ante*, who all participate at a lower intensity. This therefore is consistent with the decline in market-making observed in BIS (2014), and in various subsequent studies that are cited in footnote 2.[3]

The third moment, the bid-ask spread, declined by 6.6% per year in the course of the sample. This time series along with the ones for the first and second moment let us estimate the two remaining model parameters: c and u. We find that the cost of bidding c is relatively stable at around 0.02 basis points (growth is only 1.2% per year), and the seller's reservation value u rises over time although this trend is not statistically significant. One possible explanation is that over time end-users developed their own optimal-execution algorithms that, instead of submitting market orders that take price quotes from the book, often post price quotes themselves to "earn the spread."[4]

[3]There are several alternative explanations for the decline in the number of middlemen: the gradual decline in volatility, the decline in volume, or the gradual change in market structure where typically electronic trading has become an option for a larger set of traders. These explanations have merit, although none of them could explain more variability in liquidity supply over time. Note how the last point, the increased availability of electronic trading, might well explain why our estimate of u, the outside option of the seller, increased over time (as we argue in the next paragraph).

[4]See, for example, "Algorithmic trading: trends and existing regulation," European

Overall, the most salient feature of the estimation is a 7.8% annual decline in the participation of middlemen. This trend is consistent with the general unease about market liquidity worsening in the recent decade. Further empirical analysis beyond model estimation shows that SPY volume declined by 4.0% per year and, more alarmingly, order-book depth decreased steeply: 19.2% annually for depth at the best quote and 27.2% annually for depth within 50 basis points of the midquote.

Related literature on liquidity supply. Our findings contribute to the literature on liquidity supply in securities markets.[5] Various empirical asset-pricing studies show that agents demand higher returns for securities with either lower liquidity supply level or higher liquidity risk (e.g., Acharya and Pedersen, 2005; Pastor and Stambaugh, 2003). Theoretical studies offer various explanations for the average level of liquidity (e.g., order-processing cost, adverse-selection, or inventory risk)[6], but not for liquidity *risk*. Our model generates such risk endogenously through mixed strategies. Baruch and Glosten (2019), also generates liquidity risk through mixed strategies by modeling undercutting risk (as opposed to bidding cost which is the core friction in our model).

Our model further generates left skewness in bids, a feature that has been documented in several limit-order book studies. Biais, Hillion, and Spatt (1995, Figure 1), Goldstein and Kavajecz (2004, Figure 1), Hollifield et al. (2006, Table 3), Naes and Skjeltorp (2006, Table 2), and Degryse, de Jong, and van Kervel (2015, Table 2) document left skewness in the bid-price distribution for French, U.S., Swedish, Norwegian, and Dutch stocks, respectively.

We are not the first to study strategic bidding both theoretically and empirically. Cassola, Hortaçsu, and Kastl (2013) do the same for bidding in European Central Bank auctions for short-term funds. They find evidence consistent with strategic bidding by banks.

Related literature on bidding. We modified the endogenous bidder participation model of Levin and Smith (1994) by assuming that bidders do

Central Bank, February 13, 2019.

[5]Examples of comprehensive survey studies are O'Hara (1995), Madhavan (2000), Amihud, Mendelson, and Pedersen (2005), Vives (2010), and Foucault, Pagano, and Röell (2013).

[6]See survey studies cited in footnote 5.

not know how many others decide to bid. This removes the asymmetric equilibria in pure strategies and leaves only one symmetric equilibrium in mixed strategies.

The common value of bidders and the dispersion of equilibrium bids further places our model in the general class of models of price dispersion for a homogeneous good. Baye, Morgan, and Scholten (2006) survey these models in depth. Our equilibrium is a special case of Hausch and Li (1993) and has a close counterpart in the all-pay auction results in Baye, Kovenock, and de Vries (1996). In our model, however, participating bidders do *not* all pay their bid, but they *do* all pay their participation cost.

Another class of models is that on ask prices where the number of customers or their trading opportunities are uncertain. Prescott (1975) studies a multi-unit auction with an uncertain number of units which results in a distribution of prices that is right skewed (to the right because they are asks, not bids). Burdett and Judd (1983) similarly have an uncertain number of informed customers and also get right skewness in the distribution of ask prices.

Shilony (1977), Rosenthal (1980), and Varian (1980) feature settings in which firms have some captive customers, perhaps because of their geographical proximity or because of their lack of information about prices elsewhere. In the online appendix we convert our model to one for ask prices and compare it to Varian's model in some detail. The key difference is that in Varian's model a firm targets both uninformed (local) buyers *and* informed buyers who compare prices across firms and buy at the cheapest firm. This double-barreled objective makes the price distribution U-shaped (low prices target informed buyers, high prices target uninformed buyers). Our setting does not have "uninformed buyers."

Our contribution. In sum, we contribute to the literature in the following ways. First, we prove uniqueness of the symmetric equilibrium in mixed strategies. Second, we show that the density of all bids and of the best bid strictly increases, is convex, and bid prices therefore are negatively skewed. Third, we establish a "purification" result that establishes an isomorphism between our equilibrium in mixed strategies and the limit of equilibria in pure strategies for a model that adds dispersed private signals/values. The limit is obtained by driving private-value dispersion to zero. This should appeal to those who have reservations regarding whether equilibria in mixed strategies could speak to real-world settings because they find it hard to believe that agents would use randomization in their strategies. Fourth, we offer an empirical strategy to estimate the model parameters. As equilibrium expressions are relatively simple, they are straightforward to estimate with meaningful empirical moments. We implement the strategy on a decade of order-book snapshots and, through the lens of the model, provide an understanding of various trends in liquidity supply.

2 Model

In this section we first present the model and its unique symmetric equilibrium in mixed strategies. We then show that the equilibrium is robust in the sense that it can be obtained as a limit from a slightly altered model with equilibria in pure strategies. Finally, we discuss one application of the model that is taken to the data in the remainder of the paper.

2.1 Primitives

Consider a first-price common-value auction for an indivisible object. Its value to bidders is v and to the seller it is u. Let there be many bidders who incur an (opportunity) cost $c < v - u$ to bid. The parameters (v, c, u) are common knowledge.

Actions. A (potential) bidder chooses two actions: first whether or not to bid and then, if he decides to bid, what price to bid at. A bidder does not observe the actions of others. In particular, he does not know how many others have chosen to participate in the bidding.

Payoffs. If bidder i does not bid then he collects his reservation value which is zero. If he *does* bid and posts bid price p_i, then his payoff is $v - p_i - c$ if the bid wins and $-c$ otherwise. The payoff to the seller is $\max(u, \max_i p_i)$.

2.2 Equilibrium

We first solve the model for finitely many bidders N and then let N tend to infinity. The resulting equilibrium turns out to be attractive as it yields simple expressions in the model parameters. Moreover, the convergence appears to be fast as the equilibrium price distribution for $N = 10$ is almost indistinguishable from $N = \infty$ (see Figure 7 in Appendix A).

Proposition 1 ***(Unique symmetric equilibrium.)*** *There is a unique symmetric equilibrium to the bidding game. The number of bidders who decide to bid is Poisson distributed with mean:*

$$m = \ln \frac{v-u}{c}. \tag{1}$$

The ones participating bid by drawing their bid price from the following cumulative distribution function (CDF):

$$F(p) = 1 - \frac{1}{m} \ln \frac{v-p}{c} \tag{2}$$

on the support $[u, v-c]$*. The corresponding density is*

$$f(p) = \frac{1}{m} \frac{1}{v-p}. \tag{3}$$

The probability that no middleman shows up is:

$$P[\textit{No middleman shows up}] = e^{-m} = \frac{c}{v-u}. \tag{4}$$

The proof is in Appendix A, along with the proofs of all other propositions, lemmas, and corollaries. The proof combines elements of auction models with endogenous bidder participation going back to Harstad (1990), Hausch and Li (1993) and Baye, Kovenock, and de Vries (1996).

An important result of Proposition 1 is that the dispersion of bids is left skewed. This follows immediately from the bid-price density in (3) which strictly increases in p. This along with an even stronger property is worth stating formally as a corollary:

Corollary 1 ***(Bid-price density convex and increasing.)*** *The density of bids* $f(p)$ *strictly increases in* p*, it is convex, and prices therefore are left skewed.*

The intuition for the convex and increasing density (and therefore left skewness) is that to keep the bidder indifferent in the price support, price shading needs to get stronger for ever lower prices as the likelihood of winning declines disproportionately.

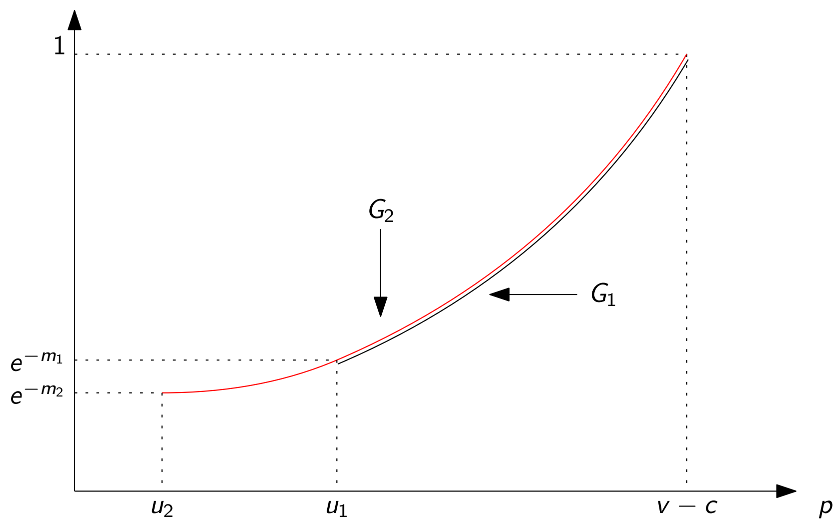

Figure 1: Best-bid distribution (by u). This figure illustrates how the distribution of the best bid G changes with the seller's reservation value u.

Corollary 2 ***(Best-bid distribution.)*** *The distribution of the best bid p is:*

$$G(p) = \begin{cases} 0 & \textit{for } p = u. \\ \frac{c}{v-p} & \textit{for } p \in (u, v-c]\,. \end{cases} \tag{5}$$

The corresponding density of the best bid strictly increases and is convex.

Figure 1 illustrates the result of Corollary 2. It plots the best-bid distributions G_1 and G_2 for model parameters that are equal except for $u_1 > u_2$. Surprisingly, they only differ in terms of their support. In other words, their values coincide for the intersection of the two supports. Changing u therefore only changes the length of the support and the size of the mass point at its lower bound. This mass point captures the probability that no middleman shows up (see (4)). The otherwise unchanged distribution, however, is surprising because the distribution of individual bids *does* depend on u (see (2)).

This surprising result follows from two economic forces that appear to exactly offset. Proposition 1 shows that when the seller is more eager to sell (lower u) then more middlemen show up on average (higher m) but each one bids less aggressively in terms of first-order stochastic dominance. The net effect is that except for the support, the best-bid distribution is invariant to the seller's reservation value.

Economically, the distribution invariance is reminiscent of Bertrand competition where a cost c is paid *ex post* (i.e., after winning the auction). Each bidder then bids $v-c$ or stays out and the seller's reservation value is therefore irrelevant for the best bid. What our model shows is that this result holds up when a cost is paid *ex ante* instead of *ex post*.

An important difference between the two cases, however, is that an *ex-ante* cost creates "crash risk." Relative to the Bertrand case, the best bid becomes stochastic with $v - c$ as an upper bound. It follows immediately from Corollary 2 that its quantiles z_α are:

$$z_\alpha = \begin{cases} u & \text{for } \alpha \in \left[0, \frac{c}{v-u}\right). \\ v - \frac{1}{\alpha}c & \text{for } \alpha \in \left[\frac{c}{v-u}, 1\right]. \end{cases} \tag{6}$$

Note that the deterministic part of price shading is the same in our model to what it is in Bertrand (i.e., $z_1 = c$). Our model, however, adds additional stochastic shading that one could say is at least as large in magnitude as the deterministic part. Its median is $\frac{1}{1/2}c - c = c$. The most biting part is the heavy left tail which for example implies that the 0.1 quantile is $10c$ and the 0.01 quantile is $100c$ (assuming $c < 0.01\,(v-u)$). This is the sense in which an *ex-ante* cost creates sizeable crash risk.

Figure 2 illustrates the crash risk implied by the model. It contains boxplots for the best-bid distribution where $v = 1$ and $u = 0$. The variation is in the cost of bidding: $c \in \{10^{-6}, 10^{-5}, 10^{-4}\}$. The low end of this range is inspired by our bidding-cost estimates which range from 0.01 to 0.10 basis points (see Section 4.2). We add the 1 basis point case to capture extreme conditions. The plots contain a box that corresponds to the interquartile range and whiskers that reach out to the 0.01 and 0.99 quantiles.

The most salient feature of the plots is the concentration of probability mass near one with a non-negligible mass far below one. More specifically, for a bidding cost of one basis point, the interquartile range is within 10 basis points of v, but the 0.01 quantile is -100 basis points. This plot in particular shows that what seems like a low cost of bidding, one basis point, could lead to a best bid that is 100 basis points below v (or, in financial terms, to a 1% value at risk (VaR) of one cent on a dollar).

2.3 Robustness

The equilibrium of Proposition 1 is robust in the sense that it can be generated after adding private signals and then taking the limit of equilibria in pure

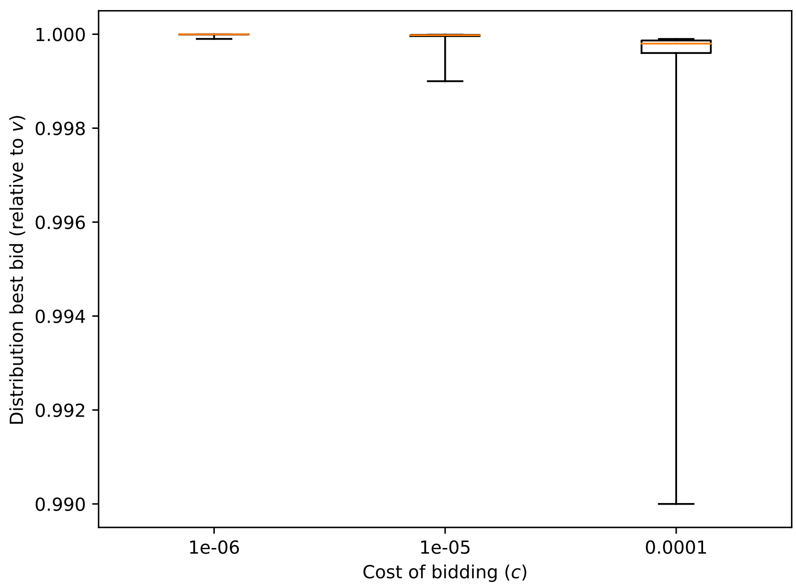

Figure 2: Best-bid distribution (by c). This graph plots the best-bid distribution through boxplots for $v = 1$, $u = 0$, $c \in \{10^{-6}, 10^{-5}, 10^{-4}\}$. The boxes depict the interquartile range and the whiskers correspond to the 1 and 99 percentiles.

strategies. Suppose that one changes the primitives in the following way. Each of the *ex-ante* identical bidders has two sources of information:

1. All bidders observe a public signal that induces a uniform common prior over $v \sim U[V-\sigma, V+\sigma]$.

2. If bidder i decides to bid and pays c, then before bidding, he observes a private signal $x_i \sim U[0, v]$.

The x_i are conditionally independent. The following proposition states the robustness result formally.

Proposition 2 ***(Purification of the Proposition 1 equilibrium.)*** *For $\sigma \downarrow 0$ the equilibrium converges to the equilibrium in mixed strategies of Proposition 1. The rate of the convergence is given by the following expression:*

$$|F_\sigma(p) - F(p)| \leq \frac{3}{c}\frac{1}{m}\left(1 - e^{-m}\right)\sigma + \mathcal{O}\left(\sigma^2\right). \tag{7}$$

The speed of convergence of F rises with the cost of participation and rises with the expected number of bidders m because the term $\frac{1}{m}\left(1 - e^{-m}\right)$ is decreasing in m. The decision whether to bid remains random; it does not depend on the private signal that is seen only if the decision to bid has been made.

2.4 Application to middlemen in limit-order markets

We believe that our model naturally applies to bidding in limit-order markets. Bidders in these markets compete for market sell orders by posting limit-order bids. When market sell orders arrive, they execute against the highest priced bids. In this sense, a limit-order market is a first-price auction.

The model parameters could be thought of as follows. The bidding cost c is the opportunity cost that a single middleman experiences when he posts a bid. He likely spends non-trivial resources to determine what price to bid at (e.g., parse recent data feeds, assess reselling opportunities). The value of winning v could be thought of as the value of reselling the position in the future (see, e.g., Grossman and Miller, 1988). The seller's reservation value could be his outside option of posting an ask price himself. This, however, is likely to be costly for him simply because he has to wait until a market buy order arrives. And while waiting, he might be adversely selected when he is not quick enough to re-price his ask on incoming news (see, e.g., Jovanovic and Menkveld, 2015).

Note that middlemen also experience an adverse-selection cost when posting their bid, but these are likely to be much smaller. The mere fact that they are middlemen allows them to amortize the technology to avoid adverse-selection, hardware and software, over many trades (see, e.g., Budish, Cramton, and Shim, 2015). A deeper understanding of how such friction affects the trading process requires a dynamic model, which is beyond the scope of our study.

3 Data

In this section we introduce the sample and present some early results on realized bid-price dispersion.

3.1 Sample

To study the dispersion of bids in real-world markets we collect snapshots of the Nasdaq limit-order book. We focus on one of the most actively traded ETFs: SPY. The database used to build the sample is: LOBSTER (lobsterdata.com). Its distinguishing feature is that it includes information on the number of shares offered at price levels beyond the best bid or ask quote. This therefore is an ideal dataset to study realized bid-price dispersions.

A drawback of picking such an active security is that the database is enormous. Downloading just a single day from LOBSTER can take more than a day. The sample spans the entire history available in LOBSTER: July 2007 through December 2018. To keep the study feasible, however, we only include every Wednesday in the first full week of each quarter thus ending up with 46 days.

We further reduce the size of these samples in two ways. First, we take snapshots of the limit-order book that are 100,000 events apart where any change to the limit-order book counts as an event (e.g., a limit-order cancellation, modification, or deletion). This keeps sample size manageable and also avoids serial dependence in snapshots. The sample we end up with contains 19,376 snapshots (approximately 19,376/46=421 per day). Second, the snapshots contain the number of shares offered at all price levels between the midquote and 50 basis points below the midquote.[7]

For completeness, let us review two important details. First, we focus only on order-book activity during regular trading hours from 9:30 until 16:00 EST. To avoid any idiosyncrasies associated with the opening or closing auction, we remove the first and last five minutes of trading and therefore only consider snapshots between 9:35 and 15:55.

Second, LOBSTER order-book snapshots only show visible orders as a trader would see them on his screen. What is excluded are fully hidden orders or the part of iceberg orders that is hidden. The fraction of bid volume that is hidden, however, is relatively low as about 90% of trading volume is generated by market orders executing against fully visible limit orders or the visible part of partially hidden limit orders.[8]

3.2 Summary statistics

Figure 3 plots standard summary statistics for the sample. A few features stand out. First, the price of SPY steadily declines from the middle of 2007 through early 2009. From then on it increases until the end of the sample in 2018. This pattern mirrors the S&P500 index since SPY is an ETF that tracks this index. It reflects the onset of the financial crisis, coined the *Great Recession*, and the recovery from it.

Second, liquidity supply shows relatively modest changes in terms of the bid-ask spread, but a strong decline in terms of depth.[9] In the course of the

[7]The raw data allow for going deeper into the book but only to a limited extent (about 200 basis points for SPY). After downloading the larger files for some random dates, we find that most of the drop in the number of bids is within the first 50 basis points and therefore decided to keep the samples manageable and pick 50 basis points as our cutoff level. We further believe that bids further out might be there for reasons other than what the model aims to capture (e.g., stub quotes to benefit from fat-finger errors).

[8]As hidden orders are not part of the model, their presence in the data, however small it may be, is likely to bias our results. If all bids were partially hidden, then there would be no bias in the estimates of the moments and therefore no bias in the parameter estimates. If, however, pure randomness makes some bids completely hidden then the bid-ask spread will be overestimated, but the estimates of the other two moments, aggressiveness and variability, will remain unbiased (because the former is based a minimum/maximum statistic whereas the other two are based on an average statistic). (13) shows that overestimation of the spread will lead to overestimation of c and therefore, by (1), underestimation of u (i.e., the bid-price support $[u, v-c]$ is shifted to the left due to the bias).

[9]Our evidence is consistent with what has been documented for the bond market. See footnote 2 for a summary.

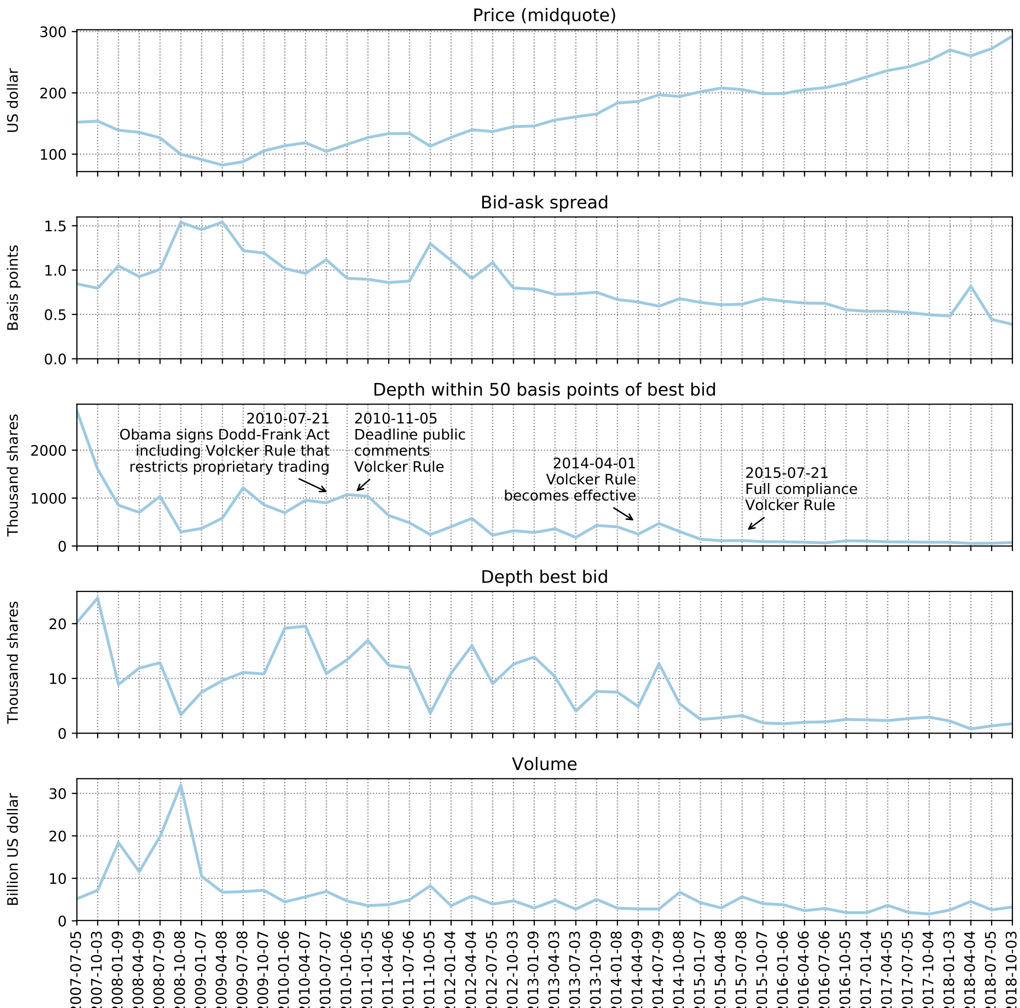

Figure 3: Summary statistics. These statistics summarize how trading in SPY evolved from 2007 through 2018. It is based on Nasdaq order-book data snapshots sampled on all Wednesdays in the first full week of each quarter.

sample, the bid-ask spread increases from slightly less than a basis point to 1.5 basis points and then steadily declines to 0.5 basis points. These changes are relatively modest as the minimum and maximum spread are only a basis point apart.

Depth, on the other hand, declines substantially, both in terms of the number of shares available at the best bid, and within 50 basis points of it. The magnitudes of these changes are striking. For 50-basis-points depth, for example, we find that at the start of the sample, it is about 2.5 million shares (at a price of around $150) dropping to about 0.5 million in the middle of the crisis and recovering to 1.0 million at the end of 2010. From that point on, it declines to less than 0.1 million at the end of 2018. The average decline from the start of 2011 until the end of the sample is 28.2% per year! The drop in best-bid depth in this period is comparable: 24.7% per year.

We believe that the depth pattern effectively implies an economically large drop in liquidity supply for large institutions starting in 2011. The reason is that such institutions often have large orders that cannot be transacted in full at the best bid price and therefore depend on depth beyond the best bid.

We conjecture that the structural break in liquidity supply is related to regulatory events in response to the crisis. On July 21, 2010, President Obama signed into law the Dodd-Frank Act which aimed to address the root causes of the financial crisis. Part of it was what came to be known as the Volcker Rule which the Board of Governors of the Federal Reserve System summarizes as follows:[10]

> Section 619 of the Dodd-Frank Wall Street Reform and Consumer Protection Act, commonly referred to as the Volcker Rule, generally prohibits insured depository institutions and any company affiliated with an insured depository institution from engaging in proprietary trading and from acquiring or retaining ownership interests in, sponsoring, or having certain relationships with a hedge fund or private equity fund. These prohibitions are subject to a number of statutory exemptions, restrictions, and definitions.

Restricting banks from proprietary trading essentially removes them as deep-pocket market makers. As is common for such far-reaching regulations, the implementation is not immediate but requires a time period for public comments, and then for potential amending the proposed regulation based

[10]See https://www.federalreserve.gov/supervisionreg/volcker-rule.htm.

on this feedback. On November 5, 2010, the request for comments closed and it is likely that by then market participants had a solid understanding of the type of regulation that would ultimately transpire. This could explain the structural break in liquidity supply for SPY by the end of 2010.

The amended Rule would become effective by April 1, 2014 with full compliance required by July 21, 2015.[11] Interestingly, the period marked by these two dates shows an accelerated drop in 50-basis-points depth. Overall, we therefore conjecture that the Volcker Rule caused a large part of the decline in depth. We will revisit the Volcker rule when discussing our result of a decline in the average number of middlemen (see Figure 6).

The third and final observation worth mentioning when parsing Figure 3 is the volume pattern. Volume shows a strong positive trend at the beginning of the sample starting at 5 billion shares per day in the middle of 2007 and peaking at around 30 billion shares per day in the fall of 2008. From then on it drops, first steeply to about 10 billion at the start of 2009, and then slowly to about 3 billion at the end of 2018.

3.3 Bid-price distributions

We close the section by examining bid-price distributions throughout the sample. Figure 4 plots the empirical densities of bid prices for each day in the sample. More specifically, it plots the *conditional* density by creating five price bins and computing how much of the total number of shares offered in these bins appears in a particular bin. The price bins are defined relative to the midquote and include: (-50, 40], . . . , (-10, 0] where all numbers are in basis points.[12]

The densities plotted in Figure 4 seem to support Corollary 1: Densities appear to be upward-sloping and convex. For all days in the sample, most bids land in the top bin. Through time, however, this top bin takes relatively more of all bids. It has approximately 50% of them at the start of the sample and 75% at the end of it. We will return to this observation after presenting the model estimates in Section 4.2 as this "slope" identifies the value of winning for middlemen.

[11] See https://www.investopedia.com/terms/v/volcker-rule.asp.

[12] The estimate of the relative mass in the right-most bin accounts for the fact that this bin is a bit shorter than the others. It runs from 10 basis points below the midquote to the best bid, not all the way through to the midquote.

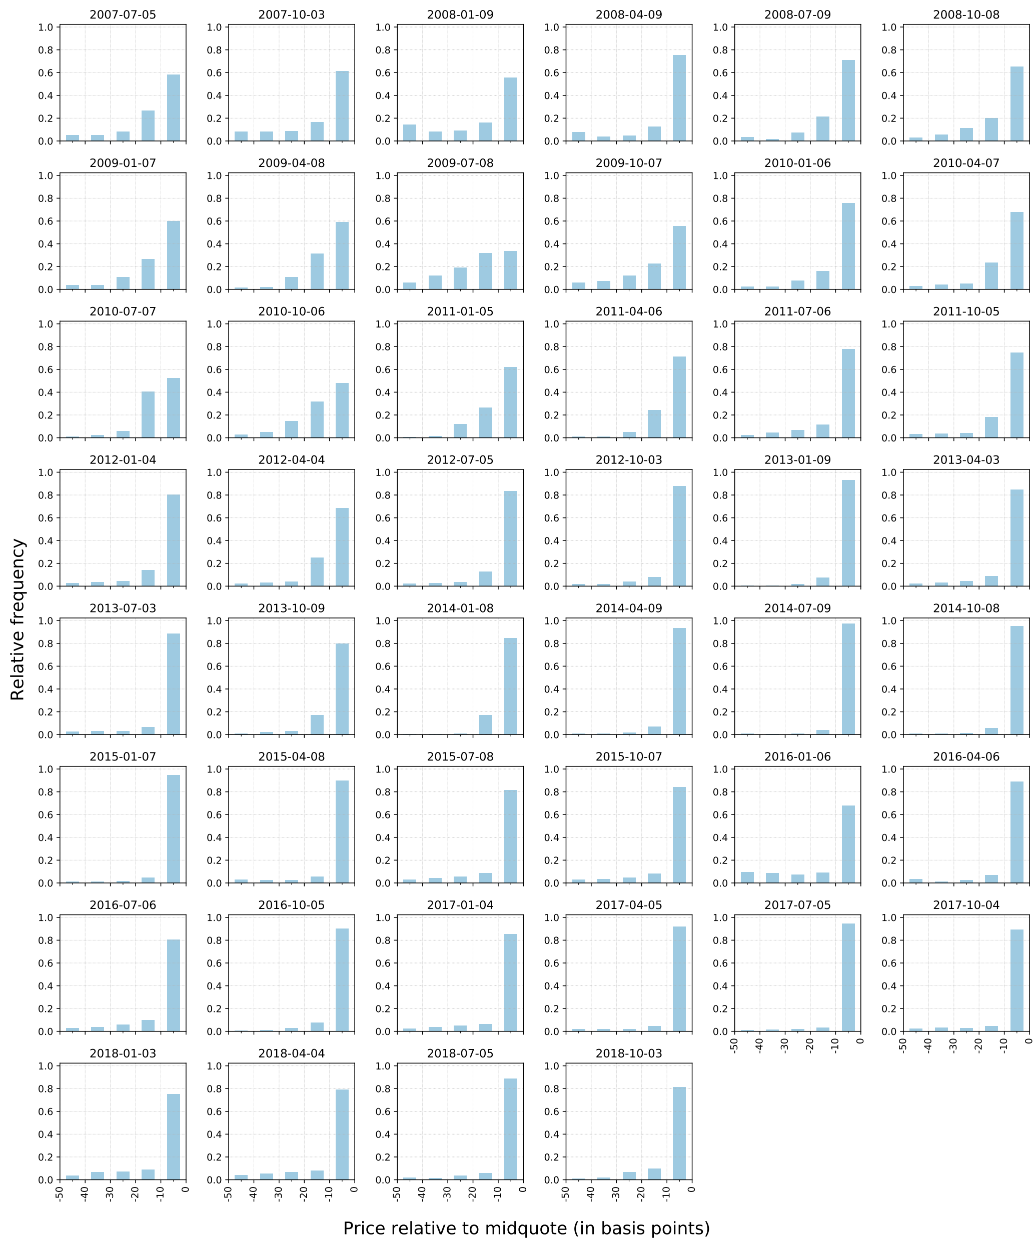

Figure 4: Bid price distribution. This figure plots the empirical densities for bids based on price bins expressed relative to the midquote: (-50 bps, -40 bps], . . . , (-10 bps, 0 bps] where "bps" is basis points. It is based on Nasdaq order-book data snapshots sampled on all Wednesdays in the first full week of each quarter.

Table 1: Testing for convexity in empirical bid-price densities. The five price bins used to estimate the bid price density allow us to compute three empirical second derivatives: $C_{i,t} = (\hat{h}_{P_i} - 2 * \hat{h}_{P_{i-1}} + \hat{h}_{P_{i-2}})/(0.001^2)$, with $i \in \{-3, -2, -1\}$ where t runs across 46 days, P_i with $i \in \{-5, -1\}$ denote price bins relative to the midquote of width 0.001, and $\hat{h}$ denotes the empirical density estimate. The first line in the table shows how many of these statistics are significantly positive thus indicating convexity. The test statistic for the full interval simply takes the average of the statistics across i. The second line in the table shows the mean of C across days divided by its standard error, which is a t statistic based on all days in the sample. $^{*}/^{**}/^{***}$ correspond to a 10/5/1 percent significance level.

	Price region relative to midquote			
	Left $C_{-3,t}$	Middle $C_{-2,t}$	Right $C_{-1,t}$	Full $\bar{C}_t$
$(\#C > 0)$/Total	29/46	41/46	43/46	45/46
$\bar{C}_i$ (in 1000)	20.1^{***}	63.3^{***}	535.1^{***}	206.2^{***}
	(0.6)	(0.8)	(1.3)	(0.3)

Although the empirical densities seem to largely agree with Corollary 1, we prefer to formally test for convexity. Bidders necessarily start shading more aggressively the further out they are from the top (possible) bid. To test this prediction, in essence, we need to verify whether the second-order derivative of the density is positive. We therefore analyze the following test statistic:

$$C_{i,t} = \frac{1}{\Delta P_i} \left(\frac{\Delta \hat{h}_t (P_i)}{\Delta P_i} - \frac{\Delta \hat{h}_t (P_{i-1})}{\Delta P_{i-1}} \right) = \frac{\hat{h}_t (P_i) - 2\hat{h}_t (P_{i-1}) + \hat{h}_t (P_{i-2})}{0.001^2} > 0, \tag{8}$$

where t, runs across days in the sample and $P_i, i \in \{-5, \ldots, -1\}$ runs across price intervals used to estimate the empirical density, with -5 denoting the price interval from 50 to 40 basis points below the midquote.

Table 1 presents the test results by price interval and across all intervals. The results largely support the hypothesized convexity. The results are strongest for the right-most price region, (-30, 0] basis points, with 43 out of 46 days rejecting the null of the second derivative being equal to zero in favor of it being positive.[13] The results for the regions (-40, 10] and (-50, 20] are 41

[13] In all tests we cluster by snapshot and thus explicitly account for non-zero correlation

out of 46 and 29 out of 46 rejections of zero in favor of a positive number, respectively. Taking all these intervals together, we find for all days except for one that the null of zero is rejected in favor of a positive number.

The table further reports tests for the full sample by assuming independence across days. It reports a sample average test statistic for each bin and for all bins combined. These statistics are all significantly positive thus providing empirical support for convexity in bid-price densities.

4 Model estimation

In this section we estimate the model based on the generalized method of moments (GMM). We express all empirical moments relative to the midquote which implies that the parameter estimates should all be interpreted relative to the midquote.[14] In the first subsection we list all moments that are matched and discuss what part of the model they identify. In the next subsection we discuss various implementation details. In the final subsection we present the estimates.

4.1 Parameter identification

Identifying the value of winning for middlemen: v. One way to identify v is by the ratio of densities at two points in the support, say $p_1 < p_2$. The model-implied ratio is:

$$\mathcal{M}_1 = \frac{f(p_2)}{f(p_1)} = \frac{v - p_1}{v - p_2}. \tag{9}$$

This value identifies v uniquely because $\mathcal{M}_1$ strictly decreases in v. What remains is to decide what prices p_1 and $p_2 > p_1$ to pick as any pair identifies v. We choose to let them correspond to the left- and right-most price bins of the empirical bid density. The benefit of choosing them as wide apart as possible is that v is most sensitive to the implied bid ratio this way. The reason is that the following partial derivative is maximized for the widest

across price bins when doing statistical inference.

[14]Note that the midquote in the model is endogenous and random because of randomness in the highest bid and the lowest ask. We assume that such randomness is small for an actively traded security like SPY and therefore ignore it in the estimation.

possible range (i.e., highest possible p_2 and lowest possible p_1):

$$\left|\frac{\partial \mathcal{M}_1}{\partial v}\right| = \frac{p_2 - p_1}{(v - p_2)^2}. \tag{10}$$

This highest possible sensitivity implies the lowest possible standard errors for the parameter estimate (as is evident from (20) that will be discussed in the estimation subsection).

Identifying the mean number of bids: m. The number of middlemen participating in the game and therefore the number of bids is Poisson distributed with mean m. For Poisson this implies that the variance is also equal to m. To compute these statistics in the data, one needs to know the size of each bid. Unfortunately, this is typically not observed because standard datasets only show the total number of shares offered at various price levels, not the individual orders that lead to this total.[15] To circumvent this issue we compute a moment that is invariant to bid size: the coefficient of variation (CV), which is defined as the standard deviation of the total number of shares offered, divided by its mean.[16] An additional benefit of using CV is that it is invariant to observing bids *only* in part of the support, which almost always is the case in limit-order data.[17] The second moment to match therefore is:

$$\mathcal{M}_2 = \frac{\sigma_x}{\mu_x} = \frac{1}{\sqrt{m}}. \tag{11}$$

The expression is quite intuitive as one would expect more variation when fewer middlemen participate on average.

[15] Sometimes a dataset does contain *flow* and *stock* data. Flow data in this case are part of the audit trail of electronic messages (e.g., limit-order submissions, cancellations, market-order submissions). Stock data are the snapshots of the order book that could, in theory, be created off of the flow data. Oftentimes, however, there are hiccups in the recording of flow data ("packet loss" in computer language) which are very costly as a single missed message will create permanent bias in subsequent order-book snapshots (see, e.g., Aquilina, Budish, and O'Neill, 2020, for an elucidating discussion). We therefore prefer not to rely on message data but base our analysis directly on the stock data that are reported by the Nasdaq itself.

[16] If the number of shares were converted to a number of middlemen by dividing by bid-order size, then this factor would appear both in the numerator and in the denominator of CV and therefore would cancel.

[17] If the probability of observing a bid in price range $[p_1, p_2]$ is α then the argument in footnote 16 applies with the factor being α.

Identifying the remaining parameters: c and u. The mean number of bids depends on the model parameters as follows (see (1)):

$$m = \ln \frac{v-u}{c}. \tag{12}$$

Given that $\mathcal{M}_2$ identifies m and $\mathcal{M}_1$ identifies v, what remains is to find one more moment to separate out the two remaining parameters: c and u. A natural candidate is the bid-ask spread. It is intuitive that such spread must relate to the value of c as the minimum spread is $2c$. The reason is that middlemen need to recover at least their cost of bidding c to participate voluntarily. It is twice c because the bid-ask spread pertains to bidding at the bid *and* ask side of the book (the latter is not explicitly modeled here but assumed to be symmetric, see Appendix A). The expected spread needs to exceed this lower bound because a middleman might lose the bidding contest and thus not recover c. It is therefore no surprise that the model-implied spread turns out to be increasing in m (see Appendix A):

$$\mathcal{M}_3 = 2cm\left(\frac{1}{1-e^{-m}}\right). \tag{13}$$

The (partial) derivative of $\mathcal{M}_3$ with respect to c turns out to be strictly positive (see (82)). The bid-ask spread therefore, for any candidate u, identifies a unique c. This ensures that the three moments jointly identify the model parameters.[18]

4.2 Estimation

The model parameters are estimated by matching the three empirical moments

$$M = \begin{pmatrix} M_1 & M_2 & M_3 \end{pmatrix}' \tag{14}$$

to their population counterparts denoted by $\mathcal{M}_\theta \in \mathbb{R}^3$ where θ contains the model parameters:

$$\theta = \begin{pmatrix} v & c & u \end{pmatrix}'. \tag{15}$$

[18] Equations (13) and (1) reveal how the model is not necessarily at odds with the stylized fact that the participation of more market makers reduces the spread (e.g., Kennedy et al., 2015, Table F.3). More middlemen and a lower spread could both be driven by a lower unobserved c.

We use the generalized method of moments (GMM) to estimate the parameters and to establish their standard errors. The model is exactly identified because the three parameters jointly determine the three moments uniquely (see discussion in the previous subsection). Straightforward algebra yields the following expressions for the estimators in terms of the empirical moments:

$$\hat{v} = \frac{M_1 p_2 - p_1}{M_1 - 1}, \tag{16}$$

$$\hat{c} = \frac{1}{2} \frac{M_3 \left(1 - e^{-M_2^{-2}}\right)}{M_2^{-2}}, \tag{17}$$

$$\hat{u} = \hat{v} - \hat{c} e^{M_2^{-2}}, \tag{18}$$

where hats denote estimators. To establish standard errors one needs the Jacobian of $\mathcal{M}_\theta$ with respect to θ:

$$J = \frac{\partial \left(\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\right)}{\partial \left(v, c, u\right)}. \tag{19}$$

The covariance matrix of the estimator $\hat{\theta}$ then equals:

$$\Sigma_{\hat{\theta}} = \left(J' \left(\mathrm{Cov}\left(M\right)\right)^{-1} J\right)^{-1}, \tag{20}$$

where $\mathrm{Cov}\left(M\right)$ is the covariance matrix of M.

4.3 Results

This section presents the estimation results. It plots them for the all days in the sample, discusses them, and closes with an observation on the likelihood of a flash crash in the sample period.

Figure 5 plots how the three empirical moments evolve in the full sample. The solid vertical lines illustrate 95% confidence intervals. They reveal that the empirical moments exhibit statistically significant trends which are therefore worthy of discussion.

First, the top plot shows that the slope in bid-price densities is positive. This is not surprising given the density plots of Figure 4. More interesting is that there is substantial time variation with peaks in the middle of the sample when bidding is particularly aggressive. The general trend seems to be one of more aggressive bidding over time as densities are steeper at the end of the sample than they are at the beginning of it.

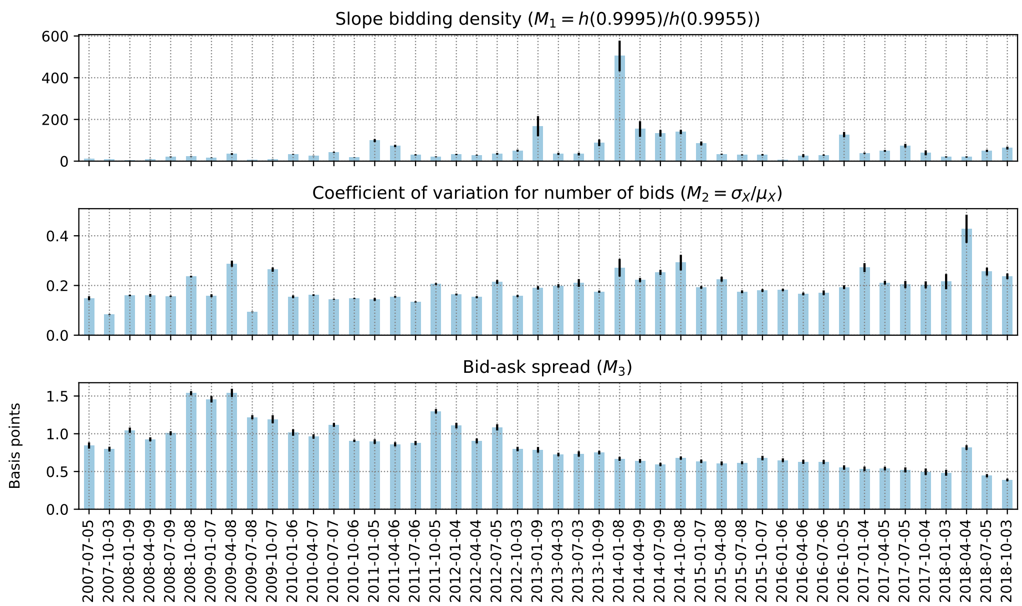

Figure 5: Empirical moments. This figure illustrates the three empirical moments that were used to estimate model parameters. It is based on Nasdaq order-book data snapshots sampled on all Wednesdays in the first full week of each quarter. Vertical solid lines denote 95% confidence intervals.

Second, the middle plot shows that the coefficient of variation hovers between 0.1 and 0.2 in the early years of the sample but gradually increases to 0.2 at the end. The overall trend therefore is more risky liquidity supply at the end of the sample.

Third, the bottom graph plots the bid-ask spread which increases in the period leading up to the financial crisis, and drops subsequently. This we already observed when discussing summary statistics.[19]

Figure 6 plots the parameter estimates implied by these empirical moments. The graphs lead to several insights. First, the top graph plots the value of winning for middleman (v). It mostly hovers between 0 and -5 basis points relative to the midquote. Winning is most valuable at the start of the sample, then drops to its lowest level in the middle of the sample and rises again towards the end of the sample. This could be due to the Volcker Rule because in the period when it came into force, 2013-2014, liquidity supply became constrained which might have meant that the value of winning became low. One reason outside of the model could be that reselling opportunities were few given that some middlemen might have been forced out of trading. If these middlemen found loopholes in the regulation or others took their place then this could explain why the value of winning increases towards the end of the sample.[20]

Second, the mean number of middlemen showing up to bid (m) exhibits a downward trend. It starts off at approximately 50 middlemen but steadily declines to approximately 20. The lion's share of this level shift seems to occur in the period when the Volcker Rule was (scheduled to be) implemented: 2013-2014. This is consistent with the Rule prohibiting some market participants from engaging in proprietary trading which includes market making. There is a slight increase in the years that follow but then a drop again towards the end of the sample.

Our finding is consistent with the predictions and the evidence on how the Volcker Rule would affect market making.[21] Prior to implementation, Duffie

[19] We note that the bid-ask spread often binds at one penny which is the tick size in this market. This is consistent with the model for a low enough c.

[20] Note that v is often below the best bid as the bid-ask spread is at most 1.5 basis points (see Figure 3). This could be explained by the model being an abstraction for *large* sell orders and middlemen with deep pockets. Large here means larger than the number of shares available at the best bid.

[21] The Volcker Rule was not a marginal rule as Deutsche Bank experienced in April 2017 when it was fined USD 157 million for violating it, along with foreign exchange rules

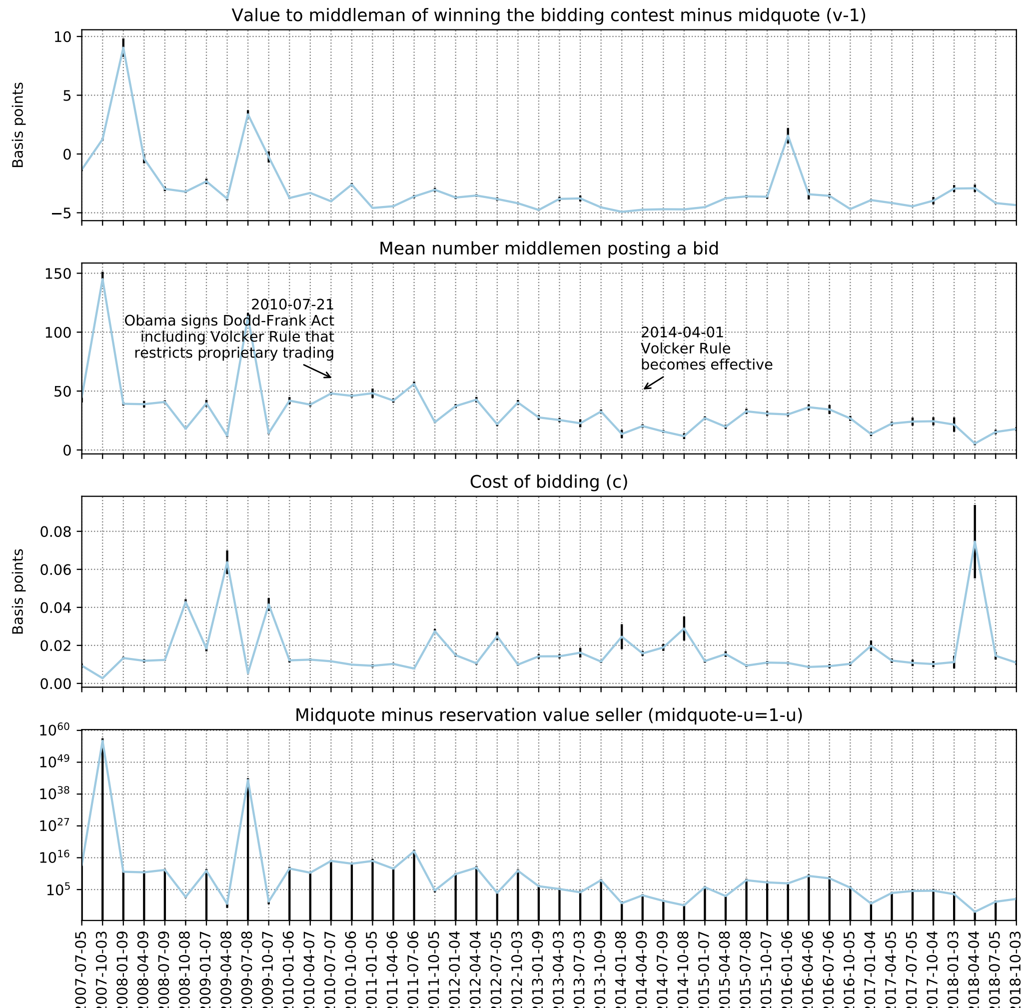

Figure 6: Parameter estimates. This figure plots parameter estimates for each day in the sample. It is based on Nasdaq order-book data snapshots sampled on all Wednesdays in the first full week of each quarter. Vertical solid lines denote 95% confidence intervals.

(2012, p. 4) writes:

> In particular, a bank that continues to offer substantial market making capacity to its clients would face a risk of regulatory sanction (and the attendant stigma) due to significant and unpredictable time variation in the proposed metrics for risk and for profit associated with changes in market prices. Likewise, the norms that are likely to arise from the proposed regulatory metrics would discourage discretion by individual market making traders in the face of career concerns. A trader's incentives and discretion would also be dampened by the proposed approach to compensation. Consequently, some banks may wish to exit the market making business.

He further observes that, in the long term, non-U.S. banks could step in to replace U.S. banks (and worries about unintended consequences of such migration outside of the regulated banking sector).

Empirical evidence on the issue is, to the best of our knowledge, only available for the bond market for which disaggregated trade is publicly available (i.e., Trade Reporting and Compliance Engine, TRACE). The most direct evidence on how the Volcker Rule affected bond liquidity is Bao, O'Hara, and Zhou (2018) who find (p. 96):

> Our second major result is that liquidity deterioration in post-Volcker stress periods featured less liquidity provision by Volcker-affected dealers, with only weak evidence of increased liquidity provision by non-Volcker-affected dealers.

This evidence is consistent with the experience of bond-market participants in the Americas. Two-thirds of them report a decrease in the number of market makers (CFA Institute, 2016). Finally, the head of the U.S. Commerce Department also mentions fewer dealers in the bond market and less market-making activity in his letter to House Committee on Financial Services (Quaadman, 2017). It confirms the predictions from a 2012 report published by his department (Thakor, 2012).

Third, the cost of bidding (c) is more or less constant in the sample period. It is approximately 0.01 basis point on average with occasional peaks of up to 0.06 basis points.

("Deutsche fined in first Volcker rule market-making case, and for forex lapses," Reuters, April 26, 2017).

Fourth, the reservation value of the seller (u) increases in the course of the sample although the increase is not statistically significant; u is estimated to be extremely low (far below zero), so much so that that we decided to plot it on a log scale. Technically, this cannot be done for negative values and we therefore decide to plot one minus u instead (i.e., the standardized midquote minus u), which is guaranteed to be positive. The bottom graph suggests that u increases over time (as one minus u decreases). Note, however, that the confidence intervals constitute a warning not to rely strongly on the level of these estimates. In most cases the interval contains zero and the hypothesis of u being zero can thus not be rejected. Mathematically, this result seems to be due to the second moment showing up exponentially in the estimator of u (see (18)). The model therefore needs really precise estimates of this second moment to produce precise estimates of u. The results show that our estimates of this moment apparently are not precise enough.

The model allows for books to be empty. The probability of such event for a Poisson distribution equals e^{-m}, where m is the Poisson mean. The average of this probability for the sample is 0.00018. If each day in the sample period were a single independent draw, then a simple calculation based on the binomial distribution suggests that the probability of not finding "empty books" in this period is 0.59. The probability, however, of finding a single empty book is 0.31. Could this have been the Flash Crash on May 6, 2010?[22]

5 Conclusion

What are the ramifications of making bidding costly for players in a first-price auction? We show that there no longer is an equilibrium in pure strategies. There is a unique symmetric equilibrium in mixed strategies. This equilibrium is easy to parse as its expressions are analytic and simple. This makes identifying the economic forces at work relatively straightforward. It further helps model estimation as standard GMM is simple to implement and parameter identification is relatively transparent. All this helps us explain the shape of bid-price dispersions in the data, why crashes happen, and why they are rare.

[22] On that day the bid-side of the book indeed ended up being virtually empty with, for some stocks, the best bid literally being a penny (SEC, 2010).

Appendices

A Proofs and derivations

Proof of Proposition 1.

The proof of Proposition 1 is done in two steps. We first prove a proposition on the equilibrium for the case of *ex-ante* finitely many middlemen. We then let this number of middlemen tend to infinity. We start by proving the following proposition.

Proposition 3 ***(Finitely many ex-ante bidders equilibrium.)*** *There are no equilibria in pure strategies. A unique symmetric Nash equilibrium exists in mixed strategies (excluding the set of correlated equilibria).*[23] *A player bids with probability*

$$\lambda = 1 - \left(\frac{c}{v-u}\right)^{\frac{1}{N-1}} \tag{21}$$

and draws a bid from the CDF

$$H(p) = \frac{1}{\lambda}\left(\frac{c}{v-p}\right)^{\frac{1}{N-1}} - \frac{1-\lambda}{\lambda}, \tag{22}$$

where H has support $[u, v-c]$. The corresponding PDF is

$$h(p) = \frac{1}{\lambda(N-1)} c^{\frac{1}{N-1}} (v-p)^{\frac{-N}{N-1}}. \tag{23}$$

The proposition is close to Hausch and Li (1993) as (21) and (22) arise in their model when their parameter α is zero and when, instead of being zero as they assume, the reservation value for the seller is u. Next, we prove each part of Proposition 3 below.

No pure-strategy equilibria. Suppose, on the contrary, that the equilibrium number of players that enter the bidding is k. Then $k = 0$ is not an equilibrium for then a sole entrant would bid $p = u$, obtain the object, and earn $v - c > 0$. Also, $k \geq 2$ is not an equilibrium for then, after the entry cost is sunk, firms would Bertrand compete on bids and all set $p = v$ and earn zero rents *ex post*, and would therefore be unable to recover the entry cost c. Finally, $k = 1$ is not an equilibrium, for then the sole bidder would bid $p = u$ and would collect a positive profit. But this would invite a second entrant who could bid $u + \varepsilon < v - c$ and, with ε small enough, would win the object and make a positive profit.

[23] If one allows for correlated equilibria then there is a trivial set of correlated equilibria whenever $N > 2$. Suppose one adds an *ex-ante* stage where M bidders with $2 \leq M < N$ are randomly selected from the set of N bidders. If then *only* these M bidders play the equilibrium identified in this proposition with $N = M$, then it is straightforward to show that such (correlated) equilibria are also equilibria, but they are asymmetric.

Proof of bid probability: (21). All bids must have an expected payoff of zero. Consider the lowest bid at $p = u$. It wins only if no other player enters, and this happens with probability $(1-\lambda)^{\text{N-1}}$. The expected payoff from such a bid is $(1-\lambda)^{\text{N-1}}(v-u) = c$, which implies (21).

Proof of bid distribution: (22). The support of H cannot be larger than $(u, v-c)$ because a bid of $p < u$ is always rejected and a bid $p > v-c$ yields a negative expected profit.

H cannot have any holes $(p_1, p_2) \subset (u, v-c)$ for then the bidder at p_2 could bid less without reducing the acceptance probability and make a positive expected profit. H cannot have mass points either (i.e., no jumps) because at any $p < v-c$ a mass point would create an incentive to shade the bid upward and a earn positive expected profit. A mass point at $p = v-c$ would imply negative profits.

All $p \in (u, v-c)$ must yield a zero expected profit net of c. This condition combined with the observation that the probability that bid p wins against k opponents with $k \sim B(N-1, \lambda)$ and with every entering bidder using H, implies:

$$\frac{c}{v-p} = \sum_{k=0}^{N-1} H^k(p) \binom{N-1}{k} \lambda^k (1-\lambda)^{N-1-k} = (\lambda H(p) + 1 - \lambda)^{N-1}. \tag{24}$$

The second equality in (24) then follows from the binomial formula. Rearranging (24) gives (22).

Convergence to Poisson for infinitely many middlemen ($N \to \infty$). Which equilibrium does one converge on when the *ex-ante* number of middlemen is taken to infinity ($N \to \infty$)? Let us first compute limit of the expected number of bids (using (21)):

$$m = \lim_{N\to\infty} N\lambda \tag{25}$$

$$= \lim_{N\to\infty} N\left(1 - \left(\frac{c}{v-u}\right)^{\frac{1}{N-1}}\right) \tag{26}$$

$$= \lim_{N\to\infty} \underbrace{\frac{N}{N-1}}_{\to 1} \underbrace{\left(1 - \left(\frac{c}{v-u}\right)^{\frac{1}{N-1}}\right) / \left(\frac{1}{N-1}\right)}_{\to \ln\left(\frac{v-u}{c}\right)} \tag{27}$$

$$= \ln\left(\frac{v-u}{c}\right) \tag{28}$$

Inserting $\lambda = m/N$ into the binomial distribution function for the number of middlemen showing up (i.e., k) yields:

$$\lim_{N\to\infty} \binom{N}{k} \left(\frac{m}{N}\right)^k \left(1 - \frac{m}{N}\right)^{N-k} \tag{29}$$

$$= \frac{m^k}{k!} \frac{N!}{(N-k)!} \frac{1}{N^k} \left(1 - \frac{m}{N}\right)^{N-k} \tag{30}$$

$$= \frac{m^k}{k!} \underbrace{\left(\frac{N}{N}\right) \cdots \left(\frac{N-k+1}{N}\right)}_{\to 1} \underbrace{\left(1 - \frac{m}{N}\right)^{-k}}_{\to 1} \underbrace{\left(1 - \frac{m}{N}\right)^{N}}_{\to e^{-m}} \tag{31}$$

$$= \frac{m^k}{k!} e^{-m}, \tag{32}$$

which is the distribution function of Poisson with mean m.

The limit of the bid-price density in (3):

$$f(p) = \lim_{N\to\infty} h(p) = \underbrace{\frac{1}{\left(1 - \left(\frac{c}{v-u}\right)^{\frac{1}{N-1}}\right) / \frac{1}{N-1}}}_{\to \frac{1}{-\ln\left(\frac{c}{v-u}\right)}} \underbrace{c^{\frac{1}{N-1}}}_{\to 1} \underbrace{(v-p)^{-\frac{N}{N-1}}}_{\frac{1}{v-p}} \tag{33}$$

$$= \frac{1}{m} \frac{1}{v-p} \tag{34}$$

and

$$F(p) = \int_u^p f(p) = 1 - \frac{1}{m} \ln\left(\frac{v-p}{c}\right). \tag{35}$$

Figure 7 illustrates how quickly the bid-price density for finite N ((23)) converges to the one for infinitely many N ((3)). The parameters used for the plot are close to the estimates for the full sample (see Section 4.3). The plot shows that for $N = 10$, the difference between the two densities becomes almost indistinguishable. This finding suggests that the smoother and more compact expressions of Proposition 1) are a reasonable approximation for bidding by ten or more middlemen as characterized by Proposition 3.

Proof of Corollary 1.

The bid-price density $f(p)$ is increasing because

$$f'(p) = \frac{1}{m} \frac{1}{(v-p)^2} > 0 \tag{36}$$

and convex because

$$f''(p) = \frac{1}{m} \frac{2}{(v-p)^3} > 0. \tag{37}$$

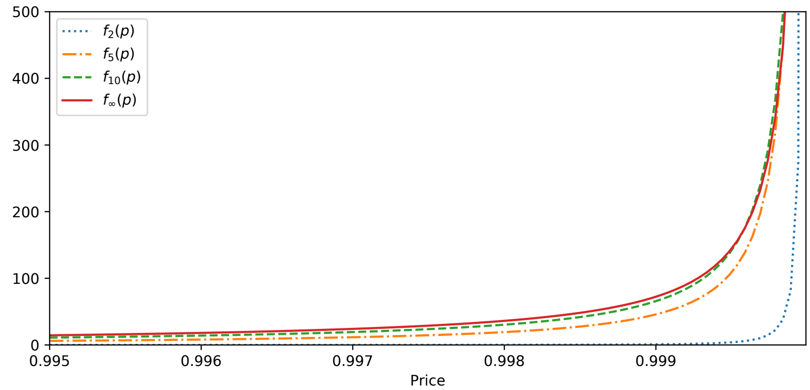

Figure 7: Speed of convergence bid-price densities. This figure illustrates how quickly bid-price densities $f_N(p)$ for the finitely many middlemen case converge to the one for infinitely many middlemen. Model parameters are: $v = 1$, $c = 10^{-6}$, and $u = 0$.

Proof of Corollary 2.

The best-bid distribution is derived by computing the conditional best-bid distribution given that k middlemen arrived, and then computing the weighted sum of these where the weights are the marginal probabilities of k:

$$G(p) = \sum_{k=0}^{\infty} F^k(p)\, P[k] \tag{38}$$

$$= \sum_{k=0}^{\infty} \left(1 - \frac{1}{m}\ln\left(\frac{v-p}{c}\right)\right)^k \frac{m^k}{k!} e^{-m} \tag{39}$$

$$= e^{-m} e^{m - \ln\left(\frac{v-p}{c}\right)} \underbrace{\sum_{k=0}^{\infty} \frac{\left(m - \ln\left(\frac{v-p}{c}\right)\right)^k}{k!} e^{-\left(m - \ln\left(\frac{v-p}{c}\right)\right)}}_{=1} \tag{40}$$

$$= \frac{c}{v-p}. \tag{41}$$

The corresponding best-bid density is $g(p) = c/(v-p)^2$, which strictly increases because

$$g'(p) = \frac{2c}{(v-p)^3} > 0 \tag{42}$$

and it is convex because

$$g''(p) = \frac{6c}{(v-p)^4} > 0. \tag{43}$$

Proof of Proposition 2

The proof of Proposition 2 is done in three steps. We first establish the zero-profit condition and the *ex-post* payoff to bid under the private signal. We then prove a lemma on the rate of convergence of the equilibrium when there are finitely many middlemen. Lastly, we let the number of middlemen go to infinity and σ go to 0.

Define the zero profit condition and payoff function. Denote the prior density by

$$\mu_\sigma(v) = \frac{1}{2\sigma} \quad \text{for} \quad v \in [V-\sigma, V+\sigma] \tag{44}$$

and the CDF and PDF of the conditionally-independent signals by, respectively,

$$\Psi(x \mid v) = \frac{x}{v} \quad \text{and} \quad \psi(x \mid v) = \frac{1}{v} \quad \text{for} \quad x \in [0, v]. \tag{45}$$

Let $\beta_\sigma(x)$ be the symmetric-equilibrium bidding strategy defined for $x \in [0, V+\sigma]$. We will drop subscript i in most of the remainder. Let the number of other bidders be $k \sim B(\lambda_\sigma, N-1)$ where λ_σ solves for λ in the zero-profit condition:

$$c = \int_{V-\sigma}^{V+\sigma} \int_0^v \pi_\sigma(x, \lambda)\, d\Psi(x \mid v)\, \mu_\sigma(v)\, dv, \tag{46}$$

where $\pi_\sigma(x, \lambda)$ is the payoff to bidding conditional on the signal x with

$$\pi_\sigma(x, \lambda) = \max_p \left(\int_0^{\beta_\sigma^{-1}(p)} (v_\sigma(x, y) - p) g_\sigma(y \mid x) dy \right), \tag{47}$$

where $y = \max_{j \neq i}(x_j)$ is the highest signal among the k other bidders, $g_\sigma(y \mid x)$ its density conditional on x, and $v_\sigma(x, y) = E(v \mid x, y)$ the conditional expectation of v given x and y. The latter two are derived below. The FOC is $(v_\sigma(x,x) - p) g_\sigma(x \mid x) dx/dp - G_\sigma(x \mid x) = 0$ which implies

$$\frac{dp}{dx} = \frac{(v_\sigma(x,x) - p) g_\sigma(x \mid x)}{G_\sigma(x \mid x)}. \tag{48}$$

Conditional density of y: $g_\sigma(y \mid x)$. Since $v \geq V - \sigma$, (45) implies that for any $x < v$, σ becomes small enough so that $x < V - \sigma$, and the support of the posterior $b(v \mid x)$ is $[V-\sigma, V+\sigma]$.[24] Conditional on (v, k), the CDF of y is $\Psi^k(y \mid v) = (y/v)^k$. Conditional on (x, k), the CDF of y is

$$G_{k,\sigma}(y \mid x) = \int_{V-\sigma}^{V+\sigma} \Psi^k(y \mid v)\, b_\sigma(v \mid x)\, dv \quad \text{for } k = 1, 2, \ldots, N-1, \tag{49}$$

[24] This result hinges on the bounded support of v which makes all derivations mathematically convenient. Importantly, it does not come at a cost of generality.

where the posterior for v is

$$b_\sigma(v \mid x) = \frac{\psi(x \mid v)\,\mu_\sigma(v)}{\int_{V-\sigma}^{V+\sigma} \psi(x \mid v')\,\mu_\sigma(v')\,dv'} = \frac{1/v}{\int_{V-\sigma}^{V+\sigma} 1/v' dv'}. \tag{50}$$

Taking the expectation of $G_{k,\sigma}(y \mid x)$ over $k \sim B(\lambda_\sigma, N-1)$ and using the binomial formula, $E\left(\Psi^k(y \mid v)\right) = (1 - \lambda + \lambda\Psi(y \mid v))^{N-1}$ with $\Psi^0(y \mid v)$ defined as having unit mass at $p = 0$, we obtain player x's predictive CDF of y as:

$$G_\sigma(y \mid x) = \int_{V-\sigma}^{V+\sigma} \left(1 - \lambda + \lambda\frac{y}{v}\right)^{N-1} b_\sigma(v \mid x)\,dv. \tag{51}$$

Its density $g(y \mid x) = \frac{d}{dy}G(y \mid x)$ is used in (47).

Conditional expectation of v: $v_\sigma(x, y)$. When calculating the expectation of v given (x, y) (i.e., $v_\sigma(x, y)$ in (47)), the bidder takes into account the winner's curse:

$$v_\sigma(x, y) = \int_{V-\sigma}^{V+\sigma} v b_\sigma(v \mid x, y)\,dv, \tag{52}$$

where the posterior density for v given (x, y) is

$$b_\sigma(v \mid x, y) = \frac{L(y \mid v)\,b_\sigma(v \mid x)}{\int_{V-\sigma}^{V+\sigma} L(y \mid v')\,b_\sigma(v' \mid x)\,dv'}. \tag{53}$$

In (53), $b(v \mid x)$ is given by (50), and the likelihood of y is

$$L(y \mid v) = \frac{d}{dy}(1 - \lambda + \lambda\Psi(y \mid v))^{N-1} \tag{54}$$

$$= (N-1)\,\lambda\psi(y \mid v)\,(1 - \lambda + \lambda\Psi(y \mid v))^{n-2}. \tag{55}$$

This completes the definition of the payoff functions *ex ante* in (46) and *ex post* in (47).

The distribution of bids. Since $\beta_\sigma(x)$ is increasing in x, a property proved in (122), the distribution of bids is

$$H_\sigma(p) \equiv \Psi\left(\beta_\sigma^{-1}(p) \mid v\right) = \frac{\beta_\sigma^{-1}(p)}{v}. \tag{56}$$

Lemma 1 ***(Rate of convergence.)*** *If $V > \sigma$, the rate of convergence, for any (fixed) $\lambda \in (0, 1)$, is*

$$|H_\sigma(p) - H(p)| \le \sigma\frac{3v - \sigma}{(v-\sigma)^2}\frac{1-\lambda}{\lambda(N-1)}\frac{v}{v-u}\left[\exp\left(\frac{\lambda(N-1)}{1-\lambda}\right) - 1\right] \tag{57}$$

for $p \in [u, V - c]$, with

$$\left|\beta_\sigma^{-1}(p) - \beta_0^{-1}(p)\right| = v\,|H_\sigma(p) - H(p)|. \tag{58}$$

The proof of Lemma 1 is voluminous and does not contain any further economic insights. We therefore moved it to the online appendix.

Proof of Proposition 2. We first show if $\sigma \downarrow 0$, then $\lambda_\sigma \to \lambda$, with λ given by (21). By definition, λ_σ solves (46). Now fix λ at an arbitrary value $\hat{\lambda} \in (0,1)$. Since $\beta_\sigma^{-1}(p)$ converges uniformly to $\beta_0^{-1}(p)$ when $\sigma \downarrow 0$ (see Lemma 1), and since $v(x,y)$ does not depend on β_σ, the function $\pi_\sigma\left(x,\hat{\lambda}\right)$ defined in (47) converges to $\pi_0\left(x,\hat{\lambda}\right)$ uniformly in x. Now $\pi_0\left(x,\hat{\lambda}\right)$ is independent of x, and therefore for all $x \in [V-\sigma, V+\sigma]$,

$$\pi_0\left(x,\hat{\lambda}\right) = \pi_0\left(0,\hat{\lambda}\right) = (v-u)\left(1-\hat{\lambda}\right)^{N-1}. \tag{59}$$

On the RHS of (59) we evaluated π_0 at the lowest signal and, hence, the lowest-x bid $\beta_0(0) = u$, which wins only if no one else bids. Then

$$\int_{V-\sigma}^{V+\sigma}\int_0^v \pi_\sigma\left(x,\hat{\lambda}\right) dH(x \mid v)\mu_\sigma(v)\, dv \overset{\sigma\downarrow 0}{\to} \pi_0\left(0,\hat{\lambda}\right). \tag{60}$$

The implicit function theorem guarantees existence of the function λ_σ in the neighborhood of $\sigma = 0$ because, using (59),

$$\frac{\partial}{\partial\lambda}\pi_0\left(0,\hat{\lambda}\right)\Big|_{\lambda=\lambda_0} = -(N-1)(v-u)\left(1-\hat{\lambda}\right)^{N-2} < 0, \tag{61}$$

where $v = V$.

From (46) this means that the RHS of (59) must equal c and therefore

$$c = (v-u)\left(1-\hat{\lambda}\right)^{N-1}. \tag{62}$$

This implies that $\lambda_\sigma \to \lambda_0$, which is the solution to (21).

At $\lambda_N = 1 - \left(\frac{c}{v-u}\right)^{1/(N-1)}$, lemma 1 states that

$$|H_\sigma(p) - H(p)| \leq \sigma\frac{3v-\sigma}{(v-\sigma)^2}\frac{1-\lambda_N}{\lambda_N(N-1)}\frac{v}{v-u}\left[\exp\left(\frac{\lambda_N(N-1)}{1-\lambda_N}\right)-1\right] \tag{63}$$

First, we take the limit as $N \to \infty$.

$$\lim_{N\to\infty}\frac{\lambda_N(N-1)}{1-\lambda_N} = \lim_{N\to\infty}\frac{1-\left(\frac{c}{v-u}\right)^{1/(N-1)}}{\left(\frac{c}{v-u}\right)^{1/(N-1)}}(N-1) \tag{64}$$

$$= \lim_{n\to\infty}\frac{\left(\frac{c}{v-u}\right)^{-1/(N-1)}-1}{(N-1)^{-1}} \tag{65}$$

Both numerator and denominator converge to zero and so we apply L'Hôpital's rule. Then

$$\lim_{N\to\infty}\frac{\lambda_N(N-1)}{1-\lambda_N} = \lim_{N\to\infty}\frac{-\left(\frac{c}{v-u}\right)^{-1/(N-1)}\ln\left(\frac{c}{v-u}\right)(N-1)^{-2}}{(N-1)^{-2}}$$
$$= -\ln\left(\frac{c}{v-u}\right) = \ln\left(\frac{v-u}{c}\right)$$

Thus

$$
\begin{aligned}
|F_\sigma(p) - F(p)| &= \lim_{N\to\infty} |H_\sigma(p) - H(p)| \\
&=\leq \sigma \frac{3v-\sigma}{(v-\sigma)^2} \ln\left(\frac{c}{v-u}\right) \frac{v}{v-u}\left[\exp\left(\ln\left(\frac{v-u}{c}\right)\right) - 1\right] \\
&= \sigma \frac{3v-\sigma}{(v-\sigma)^2} \ln\left(\frac{c}{v-u}\right)\left(\frac{v}{c} - \frac{v}{v-u}\right)
\end{aligned}
$$

Expanding the RHS around $\sigma = 0$, we get

$$
\begin{aligned}
|F_\sigma(p) - F(p)| &\leq \frac{3\sigma}{v} \ln\left(\frac{c}{v-u}\right)\left(\frac{v}{c} - \frac{v}{v-u}\right) + \mathcal{O}\left(\sigma^2\right) \\
&= \frac{3}{c} \ln\left(\frac{c}{v-u}\right)\left(1 - \frac{c}{v-u}\right)\sigma + \mathcal{O}\left(\sigma^2\right) \\
&= \frac{3}{c}\frac{1}{m}\left(1 - e^{-m}\right)\sigma + \mathcal{O}\left(\sigma^2\right)
\end{aligned}
$$

This concludes the proof of Proposition 2.

Derivation of the model-implied bid-ask spread ($\mathcal{M}_3$).

To calculate the mean bid-ask spread and derive (13), we first calculate the expected best bid, then by assuming a symmetric game on the ask side of the book we calculate the expected best ask. The expected bid-ask spread then simply is the difference between the two.

Using the distribution of the best bid $G(p)$ in (5) and applying integration by parts we find that the expected best bid is:

$$
\begin{aligned}
E(p) &= \frac{1}{1-G(u^+)} \int_{u^+}^{v-c} p dG(p) && (66) \\
&= \frac{1}{1-G(u^+)} \left(-p(1-G(p))\big|_{u^+}^{v-c} + \int_{u^+}^{v-c} (1-G(p))\, dp\right) && (67) \\
&= \frac{1}{1-G(u^+)} \left(-uG\left(u^+\right) + v - c - \int_{u^+}^{v-c} G(p)\, dp\right) && (68) \\
&= \frac{v-c-u}{1-G(u^+)} + u - \frac{1}{1-G(u^+)} \int_{u^+}^{v-c} G(p)\, dp, && (69)
\end{aligned}
$$

with

$$
1 - G\left(u^+\right) = 1 - \lim_{p\downarrow u} \frac{c}{v-p} = 1 - \frac{c}{v-u} = 1 - e^{-m} \tag{70}
$$

and

$$
\int_{u^+}^{v-c} G(p)\, dp = c\int_{u^+}^{v-c} \frac{1}{v-p} dp = -c \ln\left(\frac{v-(v-c)}{v-u}\right) = cm \tag{71}
$$

and yields

$$
E(p) = v - \frac{cm}{1-e^{-m}}. \tag{72}
$$

If one assumes the value of winning is the same on the ask side of the book, then expected best ask is:

$$E(p^a) = v + \frac{cm}{1 - e^{-m}}. \tag{73}$$

The bid-ask spread is the best ask minus the best bid and therefore equals:

$$E(\mathit{BidAskSpread}) \tag{74}$$

$$= E(p^a - p) = v + \frac{cm}{1 - e^{-m}} - \left(v - \frac{cm}{1 - e^{-m}}\right) = \frac{2cm}{1 - e^{-m}}. \tag{75}$$

The spread is increasing in c as the partial derivative is positive (see $d\mathcal{M}_3/dc$ in (82)).

B Implementation details of the GMM estimation

In this appendix we clarify how we dealt with the following two issues:

1. The covariance matrix of the empirical moments collected in M is non-trivially collected from the data. We use the Delta method (i.e., first-order Taylor expansion) to express M as a non-linear function of a larger set of moments for which the covariance matrix is trivially collected from the data.
2. The population moments are expressed in terms of $\theta = \begin{pmatrix} v & c & u \end{pmatrix}'$ along with m which is itself a non-linear function of θ. We compute the Jacobian J taking this relationship into account.

Covariance matrix of M. Let X_t consist of the following variables defined based on a snapshot of the order book at time t.

$$X_t = \begin{pmatrix} \hat{h}_t(0.9975 \times \mathit{Midquote}_t) \\ \hat{h}_t(0.9825 \times \mathit{Midquote}_t) \\ \mathit{BidDepth20}_t^2 \\ \mathit{BidDepth20}_t \\ \mathit{BidAskSpread}_t \end{pmatrix}, \tag{76}$$

where

- $\mathit{Midquote}_t$ is the average of the weighted bid and ask price where for, for example, for the bid the price is equal to depth-weighted bid prices within 20 basis points of the best bid (depth is in number of shares).
- $\hat{h}_t(0.99825 \times \mathit{Midquote}_t)$ is the number of shares available between 20 and 15 basis points from the midquote divided by the total number of shares available within 20 basis points of the midquote.
- $\hat{h}_t(0.99975 \times \mathit{Midquote}_t)$ is similarly defined but for the range 5 basis points below to the midquote to the midquote.

- *BidDepth20*$_t$ is the total number of shares available within 20 basis points of the midquote.

- *BidAskSpread*$_t$ is bid-ask spread expressed in basis points (i.e., 10,000 times 2(Ask-Bid)/(Ask+Bid)).

The mean of X_t is estimated as

$$\mu = \frac{1}{T} X_t. \tag{77}$$

and its covariance matrix as covariance matrix of X_t for the subsample under consideration is computed as:

$$Cov(X_t) = \frac{1}{T} X_t X_t' - \mu\mu'. \tag{78}$$

The empirical moments M used for estimating the parameters can be expressed in terms of μ as follows:

$$M_1 = \frac{\mu_2}{\mu_1}, \quad M_2 = \frac{\mu_3}{\mu_4^2} - 1, \quad \text{and } M_3 = \mu_5, \tag{79}$$

where μ_i is the ith element of μ. The covariance matrix of M is then obtained using the Delta method:

$$\text{Cov}(M) = J \text{Cov}(X_t) J', \tag{80}$$

where the Jacobian $J = \frac{\partial(M_1, M_2, M_3)}{\partial(\mu_1, \ldots, \mu_5)}$ is

$$J = \begin{pmatrix} -\frac{\mu_2}{\mu_1^2} & \frac{1}{\mu_1} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{\mu_4^2} & -\frac{2\mu_3}{\mu_4^3} & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{81}$$

Jacobian of $\mathcal{M}$ w.r.t. θ. In the infinitely many middlemen model the Jacobian of $\mathcal{M}$ with respect to θ becomes:

$$\frac{\partial(\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3)}{\partial(v, c, u)} = \begin{pmatrix} \frac{p_1 - p_2}{(v-p_2)^2} & 0 & 0 \\ -\frac{1}{2m\sqrt{m}(v-u)} & \frac{1}{2m\sqrt{m}c} & \frac{1}{2m\sqrt{m}(v-u)} \\ 2\left(A - A^2 m\right) & 2\left(B^2 m - B\right) & 2\left(A^2 m - A\right) \end{pmatrix}, \tag{82}$$

with

$$A = \frac{e^{-m}}{1 - e^{-m}}, \quad B = \frac{1}{1 - e^{-m}}. \tag{83}$$

Other results. Another object of interest is the expected number of middlemen that show up *ex-post* and post a bid which we define simply as:

$$\hat{m} = \ln\left(\frac{\hat{v} - \hat{u}}{\hat{c}}\right). \tag{84}$$

We use the delta method to compute its variance:

$$\text{var}(\hat{m}) = J \Sigma_{\hat{\theta}} J' \tag{85}$$

where

$$J = \frac{\partial \hat{m}}{\partial \left(\hat{v}, \hat{c}, \hat{u}\right)} = \begin{pmatrix} \frac{1}{\hat{v}-\hat{u}} & -\frac{1}{\hat{c}} & -\frac{1}{\hat{v}-\hat{u}} \end{pmatrix}. \tag{86}$$

References

Acharya, V.V. and L.H. Pedersen. 2005. "Asset Pricing with Liquidity Risk." *Journal of Financial Economics* 77:375–410.

Adrian, Tobias, Michael J. Fleming, Or Shachar, and Erik Vogt. 2017. "Market Liquidity After the Financial Crisis." *Annual Review of Financial Economics* 9:43–83.

Adrian, Tobias, Michael J. Fleming, Daniel Stackman, and Erik Vogt. 2015. "Has Liquidity Risk in the Treasury and Equity Markets Increased?" Manuscript, New York Fed blog post October 6, 2015.

Amihud, Yakov, Haim Mendelson, and Lasse Pedersen. 2005. "Liquidity and Asset Prices." *Foundations and Trends in Finance* 1:269–364.

Anderson, Mike and René M. Stulz. 2017. "Is Post-Crisis Bond Liquidity Lower?" Manuscript, NBER Working Paper No. 23317.

Aquilina, Matteo, Eric Budish, and Peter O'Neill. 2020. "Quantifying the High-Frequency Trading "Arms Race": A Simple New Methodology and Estimates." Manuscript, Financial Conduct Authority Occassional Paper 50.

Bao, Jack, Maureen O'Hara, and Xing (Alex) Zhou. 2018. "The Volcker Rule and Corporate Bond Market Making in Times of Stress." *Journal of Financial Economics* 130:95–113.

Baruch, Shmuel and Lawrence R. Glosten. 2019. "Tail Expectation and Imperfect Competition in Limit Order Book Markets." *Journal of Economic Theory* 183:661–697.

Baye, M., J. Morgan, and P. Scholten. 2006. "Information, Search, and Price Dispersion." In *Handbook of Economics and Information Systems*, edited by T. Hendershott. Amsterdam: Elsevier Press, 323–377.

Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1996. "The All-Pay Auction with Complete Information." *Economic Theory* 8:291–305.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman. 2018. "Capital Commitment and Illiquidity in Corporate Bonds." *Journal of Finance* 73:1615–1661.

Biais, B., P. Hillion, and C. Spatt. 1995. "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse." *Journal of Finance* 50:1655–1689.

BIS. 2014. "Market-making and Proprietary Trading: Industry Trends, Drivers and Policy Implications." Manuscript, Bank for International Settlements.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics* 130:1547–1621.

Burdett, Kenneth and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51:955–969.

Cassola, Nuno, Ali Hortaçsu, and Jakub Kastl. 2013. "The 2007 Subprime Market Crisis Through the Lens of European Central Bank Auctions for Short-Term Funds." *Econometrica* 81:1309—-1345.

Choi, Jaewon and Yesol Huh. 2017. "Customer Liquidity Provision: Implications for Corporate Bond Transaction Costs." Manuscript, Board of Governors of the Federal Reserve System.

Degryse, Hans, Frank de Jong, and Vincent van Kervel. 2015. "The Impact of Dark Trading and Visible Fragmentation on Market Quality." *Review of Finance (forthcoming)* .

Dick-Nielsen, Jens. 2019. "The Cost of Immediacy for Corporate Bonds." *Review of Financial Studies* 32:1–41.

Duffie, Darrell. 2012. "Market Making Under the Proposed Volcker Rule." Manuscript, Stanford.

Foucault, Thierry, Marco Pagano, and Ailsa Röell. 2013. *Market Liquidity: Theory, Evidence, and Policy.* Oxford: Oxford University Press.

Goldstein, M.A. and K.A. Kavajecz. 2004. "Trading Strategies during Circuit Breakers and Extreme Market Movements." *Journal of Financial Markets* 7:301–333.

Grossman, Sanford J. and Merton H. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43:617–633.

Harstad, Ronald. 1990. "Alternative Common-Value Auction Procedures: Revenue Comparisons with Free Entry." *Journal of Political Economy* 98:421–429.

Hausch, Donald B. and Lode Li. 1993. "A Common Value Auction Model with Endogenous Entry and Information Acquisition." *Economic Theory* 3:315–334.

Hollifield, Burton, Robert A. Miller, Patrik Sandås, and Joshua Slive. 2006. "Estimating the Gains from Trade in Limit Order Markets." *Journal of Finance* 61:2753–2804.

Hurewicz, Witold. 1958. *Lectures on Ordinary Differential Equations*. Cambridge, Massachusetts: MIT Press.

Jovanovic, Boyan and Albert J. Menkveld. 2015. "Middlemen in Limit-Order Markets." Manuscript, Vrije Universiteit Amsterdam.

Kennedy, Miles, Adam Gilbert, Aviral Rai, Nick Forrest, Manda Shemirani, James Chrispin, Yong Jing Teow, Thomas Fisher, and Robert Churcher. 2015. "Global Financial Markets Liquidity Study." Manuscript, PricewaterhouseCoopers15.

Levin, Dan and James Smith. 1994. "Equilibrium in Auctions with Entry." *American Economic Review* 84:585–599.

Madhavan, Ananth N. 2000. "Market Microstructure: A Survey." *Journal of Financial Markets* 3(3):205–258.

CFA Institute. 2016. "Secondary Corporate Bond Market Liquidity Survey Report." Manuscript, Chartered Financial Analyst (CFA) Institute.

Federal Register. 2014. "Prohibitions and Restrictions on Proprietary Trading and Certain Interests in, and Relationships with, Hedge Funds and Private Equity Funds." Manuscript, US Government Printing Office.

Naes, Randi and Johannes A. Skjeltorp. 2006. "Order Book Characteristics and the Volume-Volatility Relation: Empirical Evidence from a Limit Order Market." *Journal of Financial Markets* 9:408–432.

O'Hara, M. 1995. *Market Microstructure Theory.* Malden, MA: Blackwell Publishers.

Pastor, L. and R.F. Stambaugh. 2003. "Liquidity Risk and Expected Returns." *Journal of Political Economy* 111:642–685.

Prescott, Edward C. 1975. "Efficiency of the Natural Rate." *Journal of Political Economy* 83:1229–1236.

Quaadman, Thomas. 2017. "Examining the Impact of the Volcker Rule on Markets, Businesses, Investors, and Job Creation." Manuscript, U.S. Chamber of Commerce.

Rosenthal, Robert W. 1980. "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica* 48:1575–1579.

SEC. 2010. "Findings Regarding the Market Events of May 6, 2010." Report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues.

Shilony, Yufal. 1977. "Mixed Pricing in Oligopoly." *Journal of Economic Theory* 14:373–388.

Thakor, Arjan V. 2012. "The Economic Consequences of the Volcker Rule." Manuscript, U.S. Chamber of Commerce.

Trebbi, Francesco and Kairong Xiao. 2019. "Regulation and Market Liquidity." *Management Science* 65:1949–1968.

Varian, H.R. 1980. "A Model of Sales." *American Economic Review* 70:651–659.

Vives, Xavier. 2010. *Information and Learning in Markets: The Impact of Market Microstructure.* Princeton: Princeton University Press.