Bidding Frictions in Ascending Auctions

Aaron Barkley* **Joachim R. Groeger**† **Robert A. Miller**‡

November 22, 2019

Abstract

This paper develops an approach for identifying and estimating the distribution of valuations in ascending auctions where an indeterminate number of bidders have an unknown number of bidding opportunities. To finesse the complications for identification and estimation due to multiple equilibria, our empirical analysis is based on the fact that bidders play undominated strategies in every equilibrium. We apply the model to a monthly financial market in which local banks compete for deposit securities. This market features frequent jump bidding and winning bids well above the highest losing bid, suggesting standard empirical approaches for ascending auctions may not be suitable. We find that frictions are costly both for revenue and allocative efficiency.

JEL: C57, D44, G21

Keywords: dynamic games, auctions, jump bids, dominance, inference

*Department of Economics, University of Melbourne. Faculty of Business and Economics Building Lvl 4, 3010 VIC Australia. Email: aaron.barkley@unimelb.edu.au

†Tepper School of Business, Carnegie Mellon University. 4765 Forbes Ave, Pittsburgh, PA 15213 USA. Email: JRG@joachimgroeger.com

‡Tepper School of Business, Carnegie Mellon University. 4765 Forbes Ave, Pittsburgh, PA 15213 USA. Email: ramiller@cmu.edu

1 Introduction

This paper studies discriminatory ascending auctions where bidders face bidding frictions. Frictions arise through the random arrival of bidding opportunities, meaning that bidders are not sure when, if ever, their next bidding opportunity will arise. Many existing empirical auction applications have used the Japanese button auction model of Milgrom and Weber (1982), which has the equilibrium prediction that winning bidders pay the highest losing bidder's valuation. Frequently, data on ascending auctions contain jump bids and winning bids substantially higher than the largest losing bid, and equilibrium auction models relying on the button auction framework are not sufficiently rich to capture this bidding behavior.

In contrast, jump bidding has an intuitive interpretation in our model: a jump bid serves as insurance against being displaced by a higher offer from another bidder and not having the chance to respond due to frictions and/or inattention to the auction's progress. Each jump bid is a trade-off between leaving money on the table, in the case that the bid wins the auction and is above the bid of the highest loser, and decreasing the probability of losing due to an inability to respond to being pushed out by a higher bid. Hence our model serves to rationalize bidding behavior in ascending auctions observed across many auction platforms.

Although introducing the notion that bidders incur an opportunity cost of time to monitor bidding behavior is both intuitive and consistent with the patterns we observe in the data, solving for a perfect equilibrium within a model incorporating these features is extremely challenging. During the auction, in equilibrium a bidder recognizes her past bidding in this auction has possibly influenced her rivals' bids, and this in turn affects the terms of trade she now faces. Moreover, the manner in which her rivals respond to her earlier bids depend on their own valuations. Therefore each bidder keeps track of the information she accumulates during the auction along the equilibrium path to form more precise beliefs about the valuations of her rivals. Defining a perfect equilibrium entails her forming beliefs based on her own past deviations from equilibrium bidding behavior as well. In such models solving all the equilibria, including those with pooling or mixed strategies, is a daunting if not impossible task with current computational technology.

Yet analyzing this class of dynamic models without solving for the equilibrium characterizing the data generating process also poses serious challenges to inference on three fronts: identifying the primitives or deep parameters of the model, estimating these primitives, and comparing the current regime with counterfactuals. For example, we cannot readily point identify private valuation using the first order condition, as in a first price sealed bid auction. In an ascending auction the analogous first condition is an Euler equation in which the conditioning variables, the information about the auction history the bidder observes, depend on her valuation.

We fully specify the auction game, but only impose a subset of the conditions that define a best reply. The premise for our analysis is that bidders do not play dominated strategies, and do not bid above their valuation. Our analysis both builds upon and contrasts with the foundational work of Haile and Tamer (2003): in their incomplete specification of the auction mechanism, bidders never bid above their valuation (like our bidders) and, subject to that constraint, always increase their bid before the end of the auction if they would otherwise lose (unlike our bidders). Within our framework, it is easy to show that if there are unlimited opportunities to bid just before the auction closes, making a jump bid is a dominated strategy, and money left on the table vanishes if players do not play dominated strategies. Yet jump bids occur throughout the whole auction and there is almost invariably money left on the table; our framework rationalizes both stylized facts in an intuitive manner.

Our approach has several other attractive features. Because our assumptions on bidder behavior are weaker than imposing Bayesian rationality, our approach is robust to failures in equilibrium refinements that are not rationalizable. Although we do not assume an equilibrium induces the data generating process in each auction, let alone that all the auctions in the data are generated by the same equilibrium, our analysis is consistent with equilibria where there is jump bidding and money is left on the table. We exploit the property of our model that following undominated strategies only requires bidders to take actions that are mappings from a coarse partition of their respective information sets, rather than much more detailed refinements. Critically the partition we exploit does not include information about what bidders can infer about their rivals' valuations. Nevertheless it suffices to point identify quantiles of the bidder valuation distribution at all values of bids made during the auction that would win the auction if all bidding stopped at that point. Our identification strategy yields a simple estimator that finesses conditioning on all the state variables that determine a bidder's best response. Finally, although it is not feasible to compute the equilibrium of the game, we derive bounds on the expected value of winning valuations, and hence provide bounds on the costs of bidding frictions.

The empirical setting is the market for certificates of deposit (CDs) in the state of Texas from 2006 to 2009. The state government of Texas holds auctions up to twice a month to offer local banks the chance to bid for funds in the form of CDs. These banks then use the funds provided by the state to finance local projects. Banks compete with the interest rate they offer for funds. The auction has a discriminatory pricing rule: each winning bank pays the final, highest rate it bids. Bidding takes place over a fixed time span, and banks are free to submit as many (increasing) bids as they would like. The information available to participating banks is limited; banks are able to learn the current provisionally winning rate, but they are not able to directly observe bids submitted by other banks.[1] Each bank is able to win up to $7

[1]In previous theoretical models featuring jump bidding, including Avery (1998), Easley and Tenorio (2004), and Hörner and Sahuguet (2007), jump bids act to signal one's valuation to other bidders in order to discourage entry or future bids. As bidders in our application are unable to directly observe the actions of the other players,

million, and there are multiple winning banks in each auction. The equilibrium outcome for the button auction model in this setting predicts that all winning banks pay the same rate – the highest valuation among losing banks. However, we document several empirical patterns that contradict this equilibrium prediction and motivate our assumption of bidding frictions. Notably, there is wide dispersion in the price paid by winning banks, and we test and reject the frictionless Japanese auction model.

We formulate a model of bidding in a discriminatory price, ascending auction with randomly arriving bidding opportunities. The arrival process for bidding opportunities is governed by a set of time-specific distribution functions to allow for varying bidding intensity over the duration of the auction. Bank valuations depend on both observable and unobservable auction-specific components. The observable component is the auction reserve rate, which is benchmarked by US Treasury securities and hence represents aggregate macroeconomic conditions affecting demand for deposit funds. The unobserved component represents auction-specific factors affecting demand for deposit funds within the local markets we study. Conditional on these two auction-specific components, each bank's valuation is private and independent. Their private valuations represent the value of investment opportunities available to specific bank branches through their local lending opportunities.

For the duration of each auction, the data generating process tracks bids that are out-of-the-money (interest rates too low to win at the time they are tendered), the on-the-money interest rate (the current lowest winning bid absent further bidding) and bids that are in-the-money (that would win if no further bids are tendered), along with the banks' identities.[2] We show that the quantiles for the distribution of private bank valuations and unobserved heterogeneity are identified at every on-the-money rate generated by the probability distribution

this type of signaling is unlikely to explain the jump bidding we observe.

[2]We elaborate on these definitions in the description of the auction rules in Section 2.

over auction outcomes. Our empirical approach is based on the undominated strategy for a bank whose current highest bid is out-of-the money to bid in-the-money when given the chance if and only if the bank's valuation exceeds the on-the-money rate. The most important feature of this condition is that it depends on the current on-the-money rate, but not on anything else in the bank's information set. Using this condition for inference sidesteps accounting for how her beliefs about other banks' valuations are updated, which depend on their equilibrium behavior as reflected in the bank's full information set. Should she make an in-the-money bid, its precise value is used in identification and estimation if it later degrades to an on-the-money bid, but not for inferring her own valuation.

The distribution of bid opportunity arrival times, private values, and auction-specific unobserved heterogeneity are estimated separately for the months before and after the 2008 financial crisis. Examining the cost of bidding frictions to assess how the market for CDs was affected by the crisis, we find that a frictionless uniform price auction would have increased auction revenue by 8.54 percent in the pre-2008 financial crisis phase of sample, and by 2.63 percent in the latter periods of our sample, while the expected value of winning bids would have been up to 0.06 percentage points higher before the crisis and 0.03 higher after the crisis. A key factor driving the shrinkage between the differential outcomes of the auction mechanisms pre versus post crisis is that the probability distribution of valuations became more concentrated after the crisis, possibly reflecting the tighter banking regulations that the crisis precipitated, and this diminished the importance of which auction mechanism is used. Overall our empirical findings show bidding frictions have important allocative and revenue implications that are analytically and quantitatively comparable to the role and estimated size of liquidity costs Hollifield, Miller, and Sandås (2004) and Hollifield, Miller, Sandås, and Slive (2006) found in their structural econometric models of limit order markets.

In the remainder of this section we discuss the related literature. Section 2 describes the auction setting and our data. We develop the model in Section 3, while Sections 4 and 5 analyze identification and estimation, respectively. Section 6 contains the estimation results, and in Section 7 we use the estimates to quantify the costs of bidding frictions. Section 8 concludes.

1.1 Related Literature

This paper is related to the large empirical literature on ascending auctions. Athey and Haile (2002) and Athey and Haile (2007) discuss identification and estimation for various frictionless ascending auction models. Empirical studies include Paarsch (1997), who models an ascending auction as a button auction, Hong and Shum (2003), who estimate a parametric, affiliated values ascending auction model, and Athey, Levin, and Seira (2011), who compare equilibrium outcomes of open and sealed bid auctions. Aradillas-López, Gandhi, and Quint (2013) study identification in affiliated ascending auctions, also with the assumption of no frictions. Our model of random arrival of bids in an ascending auction is related to the theoretical model of Ambrus, Ishii, and Burns (2013). Ackerberg, Hirano, and Shahriar (2016) identify and estimate time and risk preferences in which bidders arrive randomly over time in a second price auction with information updating. We contribute to this literature by estimating bidder valuations without using an equilibrium prediction of bidding for inference. Instead, we exploit the weaker assumption that bidders use undominated strategies in order to identify and estimate the model.

Krasnokutskaya (2011) initiated a more recent literature on unobserved heterogeneity in auctions, applying the results of Kotlarski (1966) and Li and Vuong (1998) to identify the distribution of unobserved heterogeneity in first price sealed bid auctions. Much of the literature on unobserved auction heterogeneity has focused on the case of sealed bid auctions, whereas we identify the distribution of auction-specific unobserved heterogeneity in ascending auctions using only bid data.[3]

Finally, we contribute to the literature on identification and estimation of games without equilibrium restrictions on player behavior, which includes Aradillas-Lopez and Tamer (2008), Kline and Tamer (2012), and Kline (2018). This literature has primarily focused on static games, while we establish the identifying power of a restriction to undominated strategies in a dynamic auction game. Observing multiple actions from each player within the same game can bolster identification, as it becomes possible to further narrow the set of preferences that rationalize player behavior. In our setting, while each bank's final bid reveals the most information about that bank's valuation, a bank's interim bids determine the rate that other banks will need to match to remain in the auction. The dynamic evolution of winning interest rates is crucial to establishing identification as interim bids by one bank will determine the rates faced by other banks making their next bid.

2 Texas CD Auctions

The data come from monthly auctions for certificates of deposits issued by the Texas state government to local banks. These funds are issued for a fixed term of either 6 or 12 months. The total funds issued per auction are capped at either $50 million or $80 million. Banks in the auction compete for funds on the basis of the interest rate offered to the Texas government. Each bank is eligible to be awarded up to $7 million in funds. Banks may submit separate bids on up to five distinct parcels, sized in multiples of $100,000. The auction lasts 30 minutes. No limit is placed on the number of bids a bank can make during the course of the auction, but bids cannot be lowered or withdrawn once submitted, and the parcel size cannot be reduced. This

[3] Freyberger and Larsen (2017), using a similar strategy to Decarolis (2017), use the auction reserve rate as a signal for the unobserved term in ascending auctions. In our application the reserve rate is set by a national benchmark and does not fully reflect information about the local market we study.

a discriminatory ascending price auction: banks pay their highest, final bid they submitted on each of their winning parcels.

Banks see neither the identities of the other auction participants nor the bids submitted by other banks. The only information available to banks is the status of their own bids. There are three possibilities: in-the-money (INM), out-of-the-money (OUT), and on-the-money (ONM). Bids INM are provisionally winning; if the auction ends while a bid is INM then that bid wins. Bids OUT are losing bids: if the auction were to end when an OUT bid is placed, all winners would pay higher rates than the OUT bid. The ONM rate is defined as the lowest INM bid.

If two bids offer the same interest rate the auction mechanism gives priority to the first submission. This implies bids submitted at the ONM rate are not INM unless there remain unclaimed funds at the initial reserve rate. In this case a bid made at the reserve rate is both INM and ONM. Aside from that special case, the lowest bid currently INM has priority over any newly submitted bids at that rate, so banks submitting new bids must exceed the ONM rate to be INM.

While the ONM rate is not publicly listed by the auction platform, banks can make use of the information structure to infer it. This is done by submitting a sequence of incremental "creeping" bids until the ONM rate is reached. Once the ONM rate has been reached, banks can then place an INM bid if they choose. However learning the ONM rate in this way is not required, and banks may submit INM bids without first creeping to the ONM rate.

2.1 Example Auction

To illustrate how the auction mechanism works in practice, Figure 1 displays all bids in an auction held on July 20, 2008 to sell a total amount of CD funds of $50 million. The horizontal axis displays time elapsed in seconds and the vertical axis is the interest rate. The solid black line represents the ONM rate. The first bids were submitted shortly after the auction began,

Figure 1: Bid Data from Sample Auction

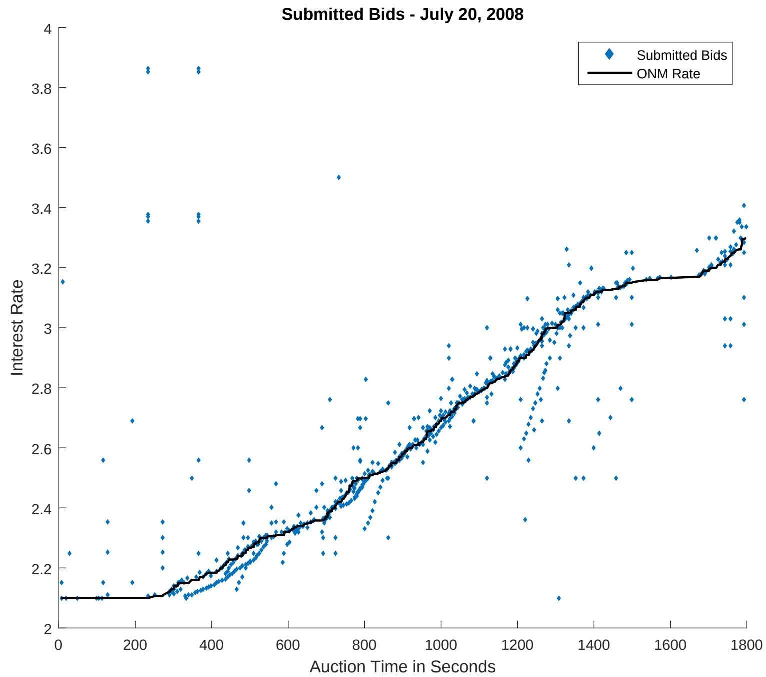

Note: The figure shows each submitted bid within one auction. Each point represents a single bid, and the solid black line is the ONM rate. The ONM rate starts at the auction reserve rate of 2.100 and increases over the course of the auction as more INM bids are placed.

and at the start of the auction all bids at or above the auction reserve rate were INM. Because each bid placed is only for a fraction of the total available funds, the ONM rate remained flat at the auction reserve rate until bids on parcels summing to $50 million had been submitted.

With reference to Figure 1, once all the funds had been committed at the reserve rate (2.100 in the example), a new INM bid of 2.106 at the 250 second mark triggered an increase to the ONM rate, displacing the most recent bid at the reserve rate of 2.100 made 230 seconds into this auction. After this initial increment to the ONM rate, each subsequent INM bid either fully or partially displaces a previous INM bid. If a bid was fully displaced, the new ONM rate was reset by the current lowest INM bid; if the new bid is for a smaller quantity than the amount bid at the current ONM rate, the rate remains unchanged. Note that new bids raising the ONM rate might themselves be partially filled. In this auction bids for $49.6 million of the available $50 million above 2.106 were placed before the auctioneer received the bid of 2.106 at 250 seconds. Hence the bank placing that bid was filled up to $400,000, but the remaining portion of the order remained OUT.

There are three notable characteristics of bidding activity in this auction. First, jump bids are common throughout the entirety of the auction's duration. Second, there is a large amount of bidding activity below the current ONM rate that constitutes incremental increases up to the ONM rate. Finally, the ONM rate steadily increases throughout the auction, continuing to push previously INM bids OUT. In the remainder of this section we detail these patterns across auctions, including the tendency for banks to leave money on that table, that is submitting winning bids at rates well above the lowest bid needed to win, the "creeping" strategy of incremental bidding used to find the ONM rate, and the reaction times to being pushed OUT.

2.2 Summary Statistics

Table 1 provides summary statistics on the 78 auctions in our sample, split into before and after the start of the 2008 financial crisis. The average number of banks in each auction pre-2008 is 28 while post-2008 the average is 23. Each bank submits an average of 14 bids per parcel, with some banks submitting up to 276 bids. Most submitted bids are provisionally winning or INM: 67 percent of bids fall in the category. The remaining bids are losing and are either ONM, meaning they are exactly at the lowest accepted winning rate but are losing due to time priority, or OUT, meaning that the bid rate is below the current ONM rate.

Most banks submit bids on one or two parcels. Changes to the size of parcels, such as increasing the desired funds on one parcel bid from $1 million to $2 million, are extremely rare. The proportion of winning banks is high: 70 percent of banks receive funds in pre-2008 auctions, and 74 percent in the post-2008 auctions.

The auction reserve rate in line (XI) is set by the current rate for a US Treasury security with the same term. This is fixed and not subject to the discretion of the state of Texas or the auction platform. All banks are aware of the reserve rate prior to the auction, and bids below the reserve rate are automatically rejected.

Table 1: Summary Statistics on Auctions

	Pre	**Post**	**Pre**	**Post**	**Pre**	**Post**	**Pre**	**Post**
	Mean		Std. dev.		Min		Max	
(I) Number of banks per auction	27.5	22.8	5.39	6.50	16.0	9.00	41.0	38.0
(II) Number of bids per parcel	13.8	12.2	20.8	23.0	1.00	1.00	186	276
(III) Proportion of bids in the money (INM)	0.67	0.69	0.13	0.12	0.31	0.44	0.87	0.87
(IV) Proportion of bids out of the money (OUT)	0.17	0.14	0.15	0.14	0.00	0.00	0.59	0.41
(V) Proportion of bids on the money (ONM)	0.17	0.17	0.07	0.06	0.08	0.08	0.42	0.32
(VI) Proportion of last bids OUT at submission	0.12	0.09	0.33	0.29	0.00	0.00	1.00	1.00
(VII) Proportion of winning bids ONM	0.08	0.07	0.07	0.05	0.00	0.01	0.27	0.19
(VIII) Size of parcels ($, in millions)	1.63	1.49	0.40	0.44	1.10	0.84	2.55	2.53
(IX) Number of parcels	1.75	1.50	1.13	0.93	1.00	1.00	5.00	5.00
(X) Proportion of banks who win	0.70	0.74	0.21	0.22	0.30	0.31	1.00	1.00
(XI) Annual reserve coupon rate	4.83	0.71	0.45	0.80	3.25	0.06	5.30	3.45
(XII) Award amount to winning bank ($, in millions)	3.88	4.51	0.48	0.57	3.20	3.57	5.08	6.31
(XIII) Jump bids ($\geq$0.01 above ONM)	0.33	0.25	0.27	0.27	0.01	0.01	3.02	1.90
(XIV) Money left on table ($):	624	1372	2116	3,607	0.00	0.00	37,660	65,380
(XV) Proportion of winning bids that were first bids	0.46	0.44	0.23	0.18	0.09	0.14	0.97	0.84
(XVI) Proportion of auctions not completely filled	0.04	0.12						

Note: The full sample consists of 24,488 bids from 78 auctions, of which 19,429 bids and 38 auctions are from pre-2008 and 5,059 bids and 40 auctions are from post-2008. We define the post-2008 period as starting September 2008 to coincide with the closure of Lehman Brothers on September 15, 2008. The total number of banks in the sample is 86, with 76 competing pre-2008 auctions and 68 competing post-2008.

Nearly all winning bids are above the minimum interest rate needed to win, and most

INM bids exceed the ONM rate by more than the minimum bid increment.[4] Jump bids more than 0.01 above the ONM rate comprise 56% of all INM bids; summary statistics on these bids are given in Row (XIII). Row (XIV) in Table 1 summarizes the money left on the table by banks that bid above the ONM rate at the end of the auction. Money left on the table is computed for each winning bid as the difference between the interest rate paid and the highest losing interest rate multiplied by the size of the parcel. The average money left on the table per winning parcel is $624 pre-2008 and $1372 post-2008; the maximum ranges $37,660 to $65,380. The prevalence of money left on the table is due to jump bids, in which banks submit bids above the lowest rate needed to be INM. Lastly, many jump bids and over 44 percent of winning bids are the first bids a bank submits.

2.3 Bidding Patterns

We now describe in greater detail some of the bidding patterns observed in the data, with an eye to documenting evidence for bidding frictions, and how banks use the information structure to learn the ONM rate. In the data we observe banks re-entering the auction and "creeping" toward the ONM rate, followed by either remaining at the ONM rate or submitting an INM rate after reaching the ONM rate.

Figure 2a plots each bid against the current ONM rate at the time of submission. Points above the 45 degree line represent bids INM, points below are OUT bids, while points on the line are ONM. (Recall that ONM bids are losing bids provided the ONM rate has gone above the auction reserve.) Increasing the ONM rate by examining points further to the right along the horizontal axis proxies for bids placed later in the auction, because the ONM rate monotonically increases throughout the auction. The graph shows all three types of bids,

[4] The minimum bid increment in the auction platform in 0.001, and 86% of all INM bids are more than this minimum increment above the ONM rate.

INM, ONM, and OUT, are common at all stages of the auction.

Figure 2b displays a plot of ONM bids against the most recent bid submitted by a bank. The tight clustering of points just below the 45-degree line shows that a preponderance of ONM bids are preceded by a slightly smaller OUT bid. This suggests that banks creep up to the ONM rate incrementally with a sequence of OUT bids.

Lastly, Figure 2c shows all INM bids that were placed just after reaching or exceeding the ONM rate. Any bid exceeding the minimal increment to the ONM rate constitutes a jump bid, as banks placing such bids are aware of the ONM rate and intentionally bid above the minimum required for INM. Over 98% of the bids in Figure 2c are greater than the minimal increment. Greater vertical distance between bids and the ONM rate indicate larger jumps; the greater dispersion from the 45-degree line in Figure 2c relative to Figure 2b is evidence that banks use a sequence of small bids to learn the ONM rate before jumping INM.

2.4 Submission and Reaction Times

This section presents evidence on the timing of bid submissions, indicating that banks do not perfectly monitor the auction. We interpret this evidence through the institutional setting: bank managers have competing uses of their time that generate opportunity costs of monitoring the auction's progress. Managing investment projects or engaging in other bank activites prevents perfect monitoring of the auction. Instead, bank managers plan their monitoring intensity before the auction begins so as to balance these other professional obligations. The chosen intensity may vary within the time span of the auction. For example, banks may want to bid more intensely at the beginning, to establish a bid at the reserve price that has high time priority, or at the end, to avoid paying too high an interest rate or being prevented from winning by a late bid from another bank.

Figure 3 presents empirical distributions over bid submission times and reaction times.

Figure 2: Bidding patterns

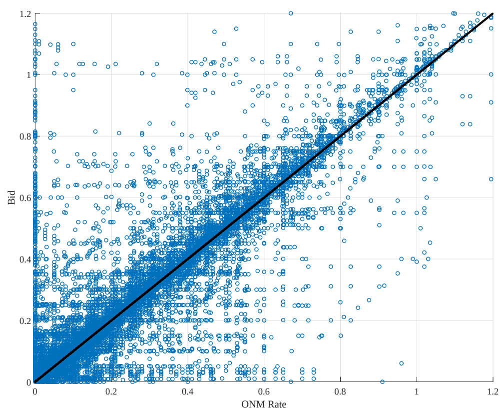

(a) All bids

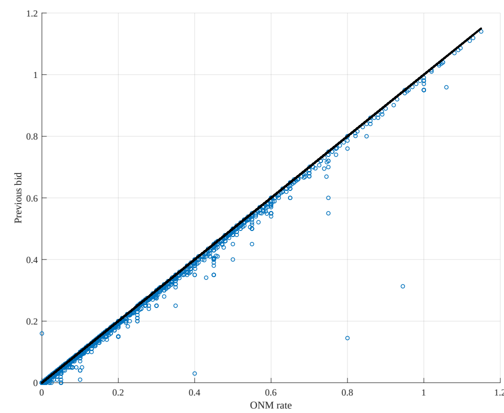

(b) Bids preceding an ONM bid

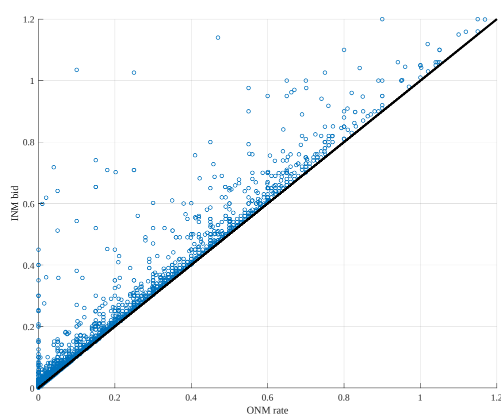

(c) INM bids made after reaching ONM

Note: These figures display bids in terms of percentage points above the reserve rate for the Texas CD auctions. Figure 2(a) shows all bids plotted against the ONM rate at the time the bid was placed. Figure 2(b) shows the bids placed just before an ONM bid was placed. Figure 2(c) plots INM that were made after reaching or exceeding the ONM rate on a previous bid and represent jumps made after reaching the ONM rate.

Figure 3a shows submission time distribution for all bids, first bids, winning bids, and last bids over the course of the 30 minute auction. More than 70 percent of first bids have been made by the 10 minute mark. The convex shape of the distribution of all bids shows that bidding activity intensifies near the end of the auction. This concentration can also be seen by looking at last bids and winning bids: more than half of all final bids are submitted in the final two minutes of the auction.

Despite the high level of activity at the end of the auction, last-minute bidding or "sniping" is not universal.[5] Figure 3a shows that 40 percent of winning bids have been submitted prior to the final minutes of the auction. Indeed some winning bids above the reserve rate are even submitted near the start of the auction. This provides evidence against a frictionless bidding environment, as banks would not submit winning bids before they had more information about rivals' valuations if they were able to perfectly monitor the auction.

We also assess how quickly banks respond to being pushed OUT to see how intensely they are monitoring the auction. Figure 3b shows empirical distributions of bank reaction times after the ONM rate has moved beyond their most recent bid. For every auction at the 5, 10, and 25 minute marks we identify all banks who have been pushed OUT in the preceding five minutes. We then measure the amount of time elapsed before a bank submits another INM bid (if ever). The distributions in the figure correspond to the quantiles of these elapsed time measurements. The distribution of reaction times at five minutes first order stochastically dominates the distributions of reaction times later in the auction, suggesting that banks monitor the auction more intensely later in the auction. However, even at late stages of the auction banks do not update the bids instantaneously upon being pushed OUT.

[5]Bajari and Hortacsu (2003) note a large tendency toward snipe bids in eBay auctions, and explain this as bidders attempting to shield information about their private signals. The information structure of the auctions in our setting makes it harder to discern other banks' actions from the observables and may help explain why sniping is less prevalent here.

Figure 3: Submission and Reaction Times

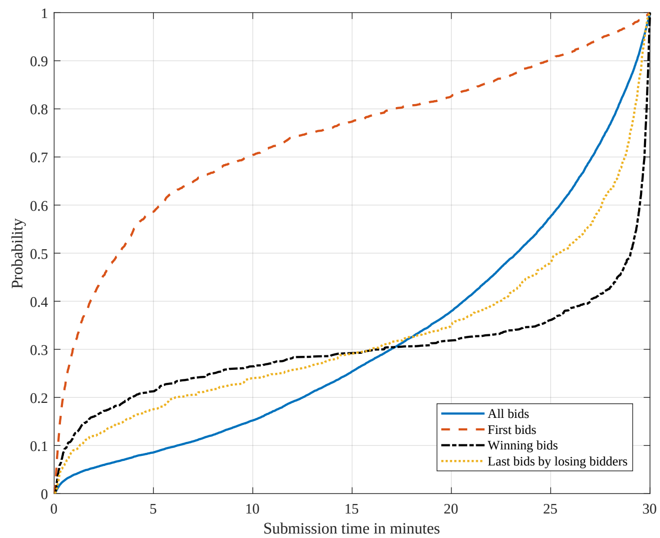

(a) Submission Times

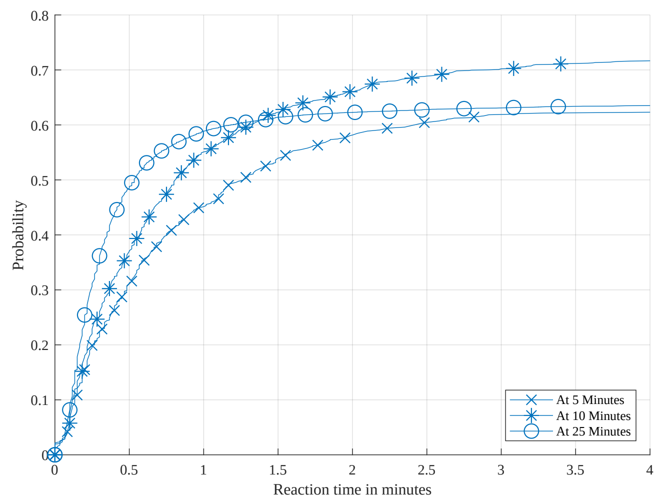

(b) Reaction Times

Note: Both figures use the full sample, combining pre- and post-2008 periods. Figure 3(a) plots the distribution of bidding times over the auction's 30 minute duration for four different bid types. Figure 3(b) shows reaction times to being pushed OUT at various stages in the auction. The probability does not necessarily increase to one as not all banks react to being pushed OUT with a bid later in the auction.

3 Model

We now develop a model of a discriminatory ascending price auction with randomly arriving bidding opportunities. The gross valuation to a typical bank from winning the auction is denoted by $\tilde{v}$ and has three components. The first is a reserve rate $\underline{r}$ set by the interest rate on a US Treasury security with the same maturity as the deposit funds issued in the auction. A second component $y \in \mathcal{Y}$ is drawn from a distribution $F_Y(\cdot)$ that is common to all banks within the auction. The last is a private valuation $x \in \mathcal{X}$ which is independently drawn by each bank from the same distribution $F_X(\cdot)$. The total gross valuation for each bank is given by:[6]

$$\tilde{v} = xy + \underline{r}$$

The private valuation component x represents the value of the funds to that particular bank, such as investment opportunities in their local area. The auction specific component y represents factors affecting the market for deposit funds in the state of Texas as a whole when the auction is held. Finally the reserve rate $\underline{r}$ captures aggregate factors affecting alternative markets for funds.

We assume the reserve rate is set exogenously, $\mathbb{E}[\ln Y] = 0$, and that X and Y are independent. Our primary concern is to identify and estimate the distribution functions $F_X(\cdot)$ and $F_Y(\cdot)$. To that end we define $v \equiv \tilde{v} - \underline{r} = xy$. We allow v to be either discrete or continuous. In both cases, we assume that the space of valuations is bounded. For v continuous we assume x and y are continuous where $\mathcal{X} \equiv [0, \overline{x}]$ and $\mathcal{Y} \equiv [0, \overline{y}]$ with $\overline{v} \equiv \overline{xy}$.

We assume each bank knows $\underline{r}$ and its own v, but not necessarily the multiplicative factors

[6] To reduce the notational burden we refrain from subscripting valuations by the identity of the bank.

x and y.[7] Bank valuations in the model are independent conditional on the auction-level component y.[8] This formulation is motivated by the fact that a bank uses funds obtained in the auction for loans and projects to the local market, and the set of projects available to one bank is unlikely to overlap much with the other banks competing in that auction. The primary factors inducing correlation of demand for funds across banks bidding in the same auction are common alternative sources of funds and general aggregate economic conditions.

Bidding occurs continuously over a time interval $\mathcal{T} = [0, T]$ with no limit on the number of bids submitted by each participant. Bidding differs from the button auction model due to the stochastic arrival of bidding opportunities. Frictions are created by a scheduling function that randomly generates bidding opportunities for banks currently OUT. Banks receiving a bidding opportunity follow a two-stage procedure, in which they first learn something about the current ONM rate and then might submit an INM bid. We assume that the process generating bidding opportunities is triggered by the start of the auction and whenever a bid is pushed OUT. The probability that a bidding opportunity is received within z seconds of being pushed OUT at time t is denoted $G_t(z)$. The bidding opportunity arrival rates vary with the time elapsed in the auction t, which allows for the overall rate of auction monitoring to be more intense at later points in the auction. Because bidding opportunities arrive randomly, the total number of times a bank can bid is a random variable. We index the order of bidding opportunities for the bank by $j \in \{1, ..., J\}$, where J, the total number of bidding opportunities, is not known until the auction is over.

The introduction of bidding frictions marks the main departure from frictionless auction

[7]This formualation is similar to Krasnokutskaya (2011), Cassola, Hortaçsu, and Kastl (2013) and several subsequent papers. A notable difference between these papers and ours is that we do not take a stand on whether the bidders know y because our estimator is robust to such changes in the information set, To the best our knowledge all previous work in this area treats the common component as known to participants and unobserved heterogeneity from an identification and estimation standpoint.

[8]Formally we assume that $E[v|v'] = E[v''|v''']$ for any four banks in the same auction with valuations v, v', v'' and v''', which implies we can define $y \equiv E[v|v']$ and $x \equiv v - y$.

models in which banks may place infinitely many bids at any point in the auction. From an empirical point of view, the assumption of no bidding frictions requires that all winning bids are relatively close to the highest losing bid. Our data on rates paid by each winning bank allows us to test this assumption directly by comparing the distribution of rates paid by the highest and lowest winning bids within an auction: under the null hypothesis of no frictions, the distributions of highest and lowest winning bids in an auction are identical. We construct a test statistic based on Li (1996) and reject this hypothesis at the 1 percent confidence level. Details of the test statistic and its asymptotic properties can be found in Appendix A.

A single bidding opportunity consists of two stages. First the bank communicates a value to the auctioneer. If that value is at least as high as the current ONM rate, the bank is told the ONM rate by the auctioneer. Otherwise the bank forfeits its opportunity to make any further bids and withdraws from the auction. The first stage mimics the creeping procedure of banks to learn the ONM rate.[9] In the second stage a bank communicating a value higher than the ONM rate in the first stage is obligated to place an INM bid. The assumption is motivated by the practical consideration that banks creep up to the ONM rate in a stepwise fashion, leading them to cross directly from OUT to INM.

Modeling bidding opportunities this way embodies two further assumptions. The first assumption is that banks who abandon bidding opportunities by falling short of the ONM rate quit the auction altogether rather than entering at some later point, while those who win the auction immediately update bids that are OUT. The evidence corroborates this assumption. An OUT bid by a losing bank is often the last bid submitted by that bank, whereas 89 percent of OUT bids submitted by winning banks are followed by another bid either within

[9] We assume that banks with sufficiently high valuations learn the ONM rate instantly. In practice, the creeping procedure used by banks takes a few seconds. The time necessary to submit bids through the online auction platform is potentially a second source of frictions in the auction; by supressing in the model we do not account for this friction. In practice, the time elapsed between subsequent bids is small relative to the time between bidding opportunities.

seven seconds or before the next increase to the ONM rate. The second assumption is that banks enter the auction at their first opportunity. This assumption is also innocuous: over 70 percent of the banks participating in an auction bid within 10 minutes. They are motivated in part by the auction format: if the sum of demand from total bids is greater than the amount supplied, but less than twice the amount supplied, only the first banks submitting at the reserve rate pay that rate, all other winners paying a higher rate.

Formally, the auction space for each bank has two components relating to the two stages of each bidding opportunity. Denote by r_j the ONM rate when the bank receives bidding opportunity j. The first stage is a value communicated to the auctioneer $c_j \in [0, \infty)$. If $c_j \geq r_j$, and only in that case, the bank learns r_j and is obligated to bid $b_j \in (r_j, \infty)$. The information set at bidding opportunity j is denoted by h_j and consists of the previous monitoring times, the corresponding ONM rates, all previous bids submitted by the bank, plus the value of y, the common component to the auction, if the bank observes it. A strategy is a mapping from the space of histories and valuations into the action space: $\beta : \mathcal{H} \times \mathcal{V} \rightarrow [0, \infty) \times (r_j, \infty)$.

Well known arguments used to establish equilibrium in private value second price sealed bid auctions prove that in this context communicating some value less than v is dominated by the strategy of selecting $c_j = v$. In practice, this means that banks creep up to the ONM rate with a sequence of incremental bids but drop out if these bids reach their own valuation before hitting the current ONM rate. Denote by W_{b_j} the event that a bid of b_j wins the auction.[10] Those with valuations higher than the current ONM rate then make an INM bid solving:

[10]Denote the minimum winning bid at the auction's conclusion by $\overline{r}$. A bid can be winning if either (i) it is strictly greater than the ONM rate at the end of the auction, $b_j > \overline{r}$, or (ii) it is equal to the final ONM rate, $b_j = \overline{r}$, and has sufficient time priority within the set of bids equal to $\overline{r}$ so that the funds awarded to bids that have higher rate or time priority is lower than the total funds amount.

$$V(h_j, r_j, v) = \sup_{b_j \in (r_j, \infty)} \left\{ \begin{array}{l} \Big[\Pr(W_{b_j}) \cdot (v - b_j) \Big] \quad + \\ \mathbb{E}\Big[\;\; \mathbb{1}\{\widetilde{W}_{b_j}\} \cdot \mathbb{1}\{j < J\} \cdot \mathbb{1}\{v > r_{j+1}\} \;\; \cdot \;\; V(h_{j+1}, r_{j+1}, v) \;\; \Big] \end{array} \right\}$$

where $\widetilde{W}_{b_j}$ is the event that the bid b_j is not winning, and J is the total number of bidding opportunities obtained by the bank. There are two cases to consider when $j = J$. The first expression arises from the current bid staying INM and the bank winning the auction. Alternatively the bid is pushed OUT and the bank never has another opportunity to bid. Juxtaposing these two outcomes incentivize the bank to make a jump bid, which reduces a bank's risk of losing the auction from not having another opportunity to bid. Analogous to a first price sealed bid auction the bank faces a trade off between making a higher bid to reduce the probability of being pushed OUT versus a lower bid that increases net value conditional on winning. The second expression models what happens when $j < J$ and the current bid is later pushed OUT but the bank obtains another chance to bid. In this case the bank only obtains positive value in the case where their valuation v and hence c_{j+1} exceeds r_{j+1}. If so, the bank has another chance to bid and solves a similar problem with an updated history.

Starting with Guerre, Perrigne, and Vuong (2000), many empirical auctions applications have used first-order conditions for optimality to derive the bid function in terms of the observables (e.g. the distribution of bids) and the model primitives. The first-order condition for the bank's problem in our model is

$$\begin{aligned} 0 \quad = \quad & \frac{\partial \Pr(W_{b_j})}{\partial b_j}(v - b_j) - \Pr(W_{b_j}) \\ & + \frac{\partial}{\partial b}\mathbb{E}\Big[\;\; \mathbb{1}\{\widetilde{W}_{b_j}\} \cdot \mathbb{1}\{j < J\} \cdot \mathbb{1}\{v > r_{j+1}\} \;\; \cdot \;\; V(h_{j+1}, r_{j+1}, v) \;\; |h_j \Big]. \end{aligned} \tag{1}$$

The first line of (1) corresponds to the first-price sealed bid case analyzed in Guerre, Perrigne, and Vuong (2000); if that was the only factor the bank's valuation can be made the subject of the resulting equation and expressed in terms of the bank's bid and the equilibrium distribution of the other bids. The second line embodies the dynamics of our model, specifically how the current bid will affect future winning probabilities through its effect on r_{j+1} and the minimum winning bid at the end of the auction. The main barrier to using (1) for inference is that both r_{j+1} and the final ONM rate depend on the other banks' strategies. These in turn depend on the (unobserved) valuation through the bank's previous bids. An ideal instrument would overcome this issue; however, any instrument making use of the history is in part a function of a bank's previous bids, which in turn depend on the unobserved valuation.[11]

4 Identification

In our Texas auction model, banks only have opportunities to bid when they are OUT, and at monitoring times they choose between exiting the auction or raising their bid to INM. If they are restricted to play undominated strategies then this discrete choice only depends on the sign of the difference between their valuation and the ONM rate. Our approach to inference leverages off two key features of our model.

First, we apply a very weak concept of equilibrium (dominance) that does not fully exploit relevant information for banks playing a more refined concept of equilibrium. For example in a Bayesian Nash equilibrium the value an INM bid takes typically depends on the whole information set of the bank and in principle could be used to infer something about the bank's valuation. The only channel though which the value of bids affect identification here is if they later become ONM due to competitive bidding, and therefore affect the decisions of

[11]The exception is the bank's first bid because the history up to that point will not depend on the bank's valuation. However, first bids provide little information about the bank's valuation.

rivals' about entering and remaining in the auction. We do not draw inferences from how an active bank sets INM after reaching ONM. This considerably simplifies the empirical analysis, transforming a computationally unmanageable problem into a tractable one.

Second, the limited way in which we use the data suffices to identify the primitives of the model up to the set of ONM rates contained in any outcome of the data generating process. In principle we sacrifice efficiency, unattainable in practice, for robustness. We do not impose the assumption that auctions are in the same equilibrium, and even allow banks to deviate from Bayesian Nash equilibrium strategies in rationalizable ways.

In principle the creeping strategy used by banks to learn the ONM rate is another source of information useful for directly inferring the valuations of the losing banks that stop creeping short of the ONM rate, presumably upon reaching their own private valuation. Although this information is straightforward to incorporate into the estimation, in practice we are concerned that the creeping might involve steps that are too big to identify bank valuations; thus our observations on creeping are used to identify an opportunity to bid, but not to pin down valuations.

4.1 A Prototype Model

To show how our approach might apply to other dynamic models, consider the following prototype of an entry-exit, commit and invest game. We define the game on the continuous time interval $t \in [0, T]$, let $\mathcal{V}$ denote the space of valuations and technology, where the valuation of a (generic) player is denoted by v is independently drawn from some unknown distribution function $F \in \mathcal{F}$ to be identified.[12] Suppose there is some probability that each player has an opportunity to move once or twice during the game. The first move occurs at some instant

[12] Following the notation of the previous section we economize on subscripting by subsuming the identity of a typical player.

$\rho \in [0, T)$, a player specific independent random variable drawn from probability distribution $G \in \mathcal{G}$. With probability g, conditional on entering the game at ρ, the player is permitted to make a (second) move at T, the end of the game; we denote by $q \in \{0, 1\}$ the outcome of whether she can ($q = 1$) or cannot ($q = 0$) make a second move. Moving entails making one or two choices on the extensive and intensive margins. The extensive margin choice, denoted by $d_t \in \{0, 1\}$ and $t \in \{\rho, T\}$, is to enter (exit) the game at ρ (at T) by setting $d_\rho = 1$ (setting $d_T = 0$) or not. The intensive margin choice, denoted by $b_t \in [0, v)$, is conditional on entering (not exiting) the game, and involves investing resources with returns that depend on the moves of everybody. Each player knows her own valuation v but not those of her rivals. When the player moves at $t \in \{\rho, T\}$ she is privy to some features of the game history, denoted by $h_t \in \mathcal{H}_t$ that include private information.

If the player does not enter at ρ (denoted by $d_\rho = 0$), she receives a zero payoff; this is also true if she enters ($d_\rho = 1$), gets the opportunity to reconsider ($q = 1$) and exits ($d_T = 0$). If she enters ($d_\rho = 1$) and does not have the chance to reconsider ($q = 0$) she receives an (expected) reward of $R_1(v, \rho, h_\rho, b_\rho)$. Finally if she enters ($d_\rho = 1$), and when given the opportunity ($q = 1$) revises her investment ($d_T = 1$) rather than exiting, she receives $R_2(v, h_T, b_\rho, b_T)$. Denote by:

$$(D_1(v, \rho, h_\rho), B_1(v, \rho, h_\rho)) : \mathcal{V} \times [0, T) \times \mathcal{H}_\rho \to \{0, 1\} \times [0, v]$$

the player's strategy at ρ; similarly define $D_2(v, h_T, b_\rho)$ and $B_2(v, h_T, b_\rho)$ for those with an opportunity to move at T.[13]

[13] Our application differs from this stylized abstraction in several key respects. In the Texas auction model more than two bids are permitted. We also identify and estimate G in a first stage, given the assumption that $G_\tau$is independent of v, and explain below how to proceed if the monitoring technology is conditional on the valuation draw.

Suppose the data are formed from observations on:

$$x_n \equiv (\rho_n, h_{n\rho}, d_{n\rho}, d_{n\rho}b_{n\rho}, h_{nT}, q_n d_{n\rho} d_{nT}, q_n d_{n\rho} d_{nT} b_{nT}) \in \mathcal{X}$$

for players $n \in \{1, \ldots, N\}$ randomly sampled within and also across different games, where h_{nt} is the (incomplete) history n knows about her own game at $t \in \{\rho_n, T\}$; thus $q_n d_{n\rho} d_{nT} b_{nT}$ captures the idea that b_{nT} is only observed if n enters the game at ρ_n and does not exit at T, and so forth. We seek to identify $F \in \mathcal{F}$, the unknown true probability distribution of preferences generating x_n. To minimize the notational burden we assume for the moment that G is known.

Consider the challenge of exploiting an interior first order condition for b_{nT} that players solve when making a final bid. Here the payoff $R_2(v_n, h_{nT}, b_{n\rho}, b_{nT})$ depends on h_{nT}, which typically depends on $b_{n\rho}$ (because rivals of n react to her first move), greatly complicating the problem of solving for v_n without explicitly computing the equilibrium at each candidate parameterization in $\mathcal{F}$. Focusing on the first order condition for the first bid at ρ_n avoids this problem because $h_{n\rho}$ does not depend on v_n. However, it introduces another problem in dealing with the derivative of the continuation value in the Euler equation, which arises because investment revisions at T depend on h_{nT}, and hence $b_{n\rho}$ and therefore v_n.[14]

Although the best response is typically a complicated function of the entire state space to be solved simultaneously for all players, focusing on some parts of the best response may be more manageable. For example suppose that applying the dominance principle is sufficient to derive the discrete choice components of the decision $D_1(v_n, \rho_n, h_{n\rho})$ and $D_2(v_n, h_{nT}, b_{n\rho})$. For each x_n we define a correspondence $\psi(x_n) : \mathcal{X} \to \Sigma_{\mathcal{V}}$, where $\Sigma_{\mathcal{V}}$ is the Borel σ-algebra generated by $\mathcal{V}$, which picks out Borel sets of $\mathcal{V}$ from sampling x_n. The sets created by $\psi(x_n)$

[14]In our application where players can move more than twice and essentially no one moves at the very last instant of the auction, all moves after the first are complicated by both factors.

form a partition of $\mathcal{V}$ that are weighted to reflect the underlying primitive F, as filtered by the known decision rules $D_1(v_n, \rho_n, h_{n\rho})$ and $D_2(v_n, h_{nT}, b_{n\rho})$, the timing of the first move stochastically determined by G, and if she entered at ρ the chance to move a second time given by the probability g. Thus F is identified if the elements of $\Sigma_{\mathcal{V}}$ coincide with the σ-algebra formed from sampling $\mathcal{X}$ and applying $\psi(x_n)$. Intuitively, for any $\widehat{v} \in [\underline{v}, \overline{v}]$, the event that a valuation is less than or equal to $\widehat{v}$ corresponds to events observed in the data that exploit known properties of the decision rules.

In the two-move analogue to our empirical application r_t is the ONM rate, $h_{n\rho} = (\rho_n, r_{n\rho})$ and $h_{nT} = (h_{n\rho}, d_{n\rho} b_{n\rho}, d_{n\rho} r_{nT})$. The equilibrium entry/exit decision for n at $t \in \{\rho_n, T\}$ is given by $d_{nt} = 1$ if and only if $v_n \geq r_t$, while the equilibrium bid of n at T depends on (v_n, h_{nT}). The entry/exit decisions are implied by the dominance principle, whereas the equilibrium bid is a complicated function that can only be derived in conjunction with a complete characterization of the equilibrium, set valued if there are multiple equilibria. The σ-algebra formed from $\psi(x_n)$ for $n \in \{1, \ldots, N\}$ and $t \in \{t_n, T\}$ is constructed from sets of the form:

$$\psi(x_n) = \begin{cases} \{v_n < r_{n\rho}\} & \text{for } \{d_{n\rho} = 0\} \\ \{v_n \geq r_{n\rho}\} & \text{for } \{d_{n\rho} = 1\} \cap \{q_n = 0\} \\ \{v_n \geq r_{nT}\} & \text{for } \{d_{n\rho} = 1\} \cap \{q_n = 1\} \cap \{d_{nT} = 1\} \\ \{r_{nT} > v_n \geq r_{n\rho}\} & \text{for } \{d_{n\rho} = 1\} \cap \{q_n = 1\} \cap \{d_{nT} = 0\} \end{cases}$$

where we abbreviate $\{x_n : d_{n\rho} = 0\}$ with $\{d_{n\rho} = 0\}$ and so forth.

4.2 Identifying $F_V(\cdot)$

Identification in this model is complicated by the existence of auction specific effects. First we analyze the identification of $F_V(\cdot)$, the probability distribution function for $V \equiv XY$, the valuation net of the reserve rate. This leads to the identification of the joint distribution of two valuations from the same auction, which we denote by $F_{V_1,V_2}(\cdot)$. Appealing to results from Kotlarski (1966) identifying $F_X(\cdot)$, the distribution of private values, and $F_Y(\cdot)$, the distribution of auction specific components, then follows.

We index a (given) bank's bidding opportunities within a (given) auction by $j \in \{1, ..., J\}$, the ONM rate increments by $s \in \{1, ..., S\}$, and their corresponding times by t_s, where $t_{S+1} \equiv T$. Let $\rho_j \in [0, T]$ denote the time the bank makes its j^{th} INM bid; let $\tau_j \in [\rho_j, T]$ denote the j^{th} time the bank is pushed OUT. Recall $r_{\rho j}$ is the ONM rate at the bank's j^{th} bidding opportunity; we define $r_{\tau j}$ as the ONM rate that subsequently pushes j^{th} bid OUT at τ_j. Define $S_J \in \{1, ..., S\}$ so that $t_{S_J} = \tau_J$, thus indexing the J^{th} time the bank is pushed OUT. Having revealed its valuation is at least r_{ρ_J} at time ρ_J, the probability a bank never reenters the auction after being pushed out at time τ_J is:

$$H(F_V, G_{\tau_J}) \equiv (1 - G_{\tau_J}(T)) + \sum_{s=S_J}^{S} \left[G_{\tau_J}(t_{s+1}) - G_{\tau_J}(t_s)\right]\left[F_V(r_s) - F_V(r_{\rho_J})\right]. \tag{2}$$

The first term, $1 - G_{\tau_J}(T)$, is the probability that another bidding opportunity is never received, preventing the bank from responding with another bid. Each component in the second term of the summation represents the joint probability that a bidding opportunity is received between the ONM rate increments and the bank's valuation lies between the ONM rate when it last bid and the current ONM rate when the opportunity is received: if a bidding opportunity is reached between t_s and t_{s+1}, which occurs with probability $G_{\tau_J}(t_{s+1}) - G_{\tau_J}(t_s)$, then if the bank does not bid it must have been due to the ONM rate passing the bank's valuation, implying that $r_{\rho_J} < v \le r_s$. Taking the sum over all increments to the ONM rate yields the total probability that the bank did receive another chance to bid but did not submit an INM bid because the required rate is too high.

The likelihood that a bank pushed OUT at τ_{j-1} submits an INM bid at time ρ_j is obtained by taking the density associated with receiving a bid at ρ_j, multiplied by the probability the bank's valuation is at least as high as r_{ρ_j} conditional on making an INM bid at ρ_{j-1}:

$$g_{\tau_{j-1}}(\rho_j) \times \Pr(v > r_{\rho_j} | v > r_{\rho_{j-1}}) = g_{\tau_{j-1}}(\rho_j) \left[\frac{1 - F_V(r_{\rho_j})}{1 - F_V(r_{\rho_{j-1}})} \right] \tag{3}$$

Because later bids contain strictly more information about a bank's valuation than earlier bids, the bound on a bank's valuation created by a new bid is a strict subset of the bounds created by all the earlier bids. This implies the likelihood pertaining to the valuations does not depend on any bids preceding the J^{th} (the last). Equation $(C.1)$ of Appendix C.1 proves the contribution of the bank to the overall likelihood is:

$$\prod_{j=1}^{J} g_{\tau_{j-1}}(\rho_j) \times \Big([1 - F_V(r_{\rho_J})] + \mathbb{1}\{\widetilde{W}_{b_J}\} H(F_V, G_{\tau_J}) \Big). \tag{4}$$

where $\mathbb{1}\{\widetilde{W}_{b_J}\}$ indicates that b_J was not a winning bid, i.e. that the bank lost the auction.

The contribution to the likelihood in equation (4) has two components. The first is a lower bound on the bank's valuation implied by the last bid, which is given by $1 - F_V(r_{\rho_J})$. This lower bound applies to all final bids, whether it is a winning bid or losing bid. The second component, $H(F_V, G_{\tau_J})$ provides upper bounds in the event that a bid is losing.

We establish identification of the quantile distribution of $F_V(\cdot)$ for every partition of the valuation space induced by the realized ONM rates. Appendix C. 3 proves the likelihood (4) is concave in the quantiles of $F_V(\cdot)$, thus establishing $F_V(\cdot)$ itself is identified at all the ONM rates that are outcomes of the data generating process. It is straightforward to show that every INM bid that is an outcome of the data generating process is also an ONM outcome in some auction because there is strictly positive probability of enough banks with high valuations to bid above every INM bid below $\overline{v}$. Thus $F_V(\cdot)$ is identified at all INM bids.

4.3 Identifying $F_X(\cdot)$ and $F_Y(\cdot)$

The probability distributions of the private signal $F_X(\cdot)$ and the auction-specific component $F_Y(\cdot)$ are identified by decomposing $F_V(\cdot)$ using the joint distribution of two valuations draws from the same auction, denoted by $F_{V_1,V_2}(\cdot)$. The identification of $F_{V_1,V_2}(\cdot)$ follows the same arguments for $F_V(\cdot)$; the joint distribution is also identified at every quantile corresponding to an ONM rate outcome of the data generating process, because its likelihood is strictly concave.

We apply deconvolution methods to a log transformation of the joint distribution (V_1, V_2) to identify $F_X(\cdot)$ and $F_Y(\cdot)$. Denote the characteristic function of the joint distribution $F_{\ln V_1,\ln V_2}(\cdot)$ by $\phi_{\ln V_1,\ln V_2}(t_1,t_2)$, let $\phi_{\ln X}(t)$ denote the characteristic function of the log private signal distribution $F_{\ln X}(\cdot)$, and $\phi_{\ln Y}(t)$ the characteristic function of the log auction-specific component $F_{\ln Y}(\cdot)$. Recalling that $V = XY$, we have that $\ln V$ is additive in $\ln X$ and $\ln Y$. Noting that X is independent across banks, it now follows from the lemma of Kotlarski (1966) that:

$$\phi_{\ln V_1,\ln V_2}(t_1,t_2) = \phi_{\ln Y}(t_1+t_2)\phi_{\ln X}(t_1)\phi_{\ln X}(t_2).$$

The expressions for the characteristic functions of $\ln Y$ and $\ln X$ are derived from this equation, with the formulation below coming from equations 2.64 and 2.65 of Rao (1992):

$$\phi_{\ln Y}(t) = \exp\left(\int_0^t \frac{\partial}{\partial u}\left[\frac{\phi_{\ln V_1,\ln V_2}(u,v)}{\phi_{\ln V_1,\ln V_2}(u,0)\phi_{V_1,V_2}(0,v)}\right]_{u=0} dv + it\mathbb{E}[\ln Y]\right) \quad (5)$$

$$\phi_{\ln X}(t) = \frac{\phi_{\ln V_1,\ln V_2}(t,0)}{\phi_{\ln Y}(t)} \quad (6)$$

Since the natural logarithm is a monotone transformation, these two equations uniquely determine the distributions of X and Y up to the location assumption $\mathbb{E}[\ln Y] = 0$.[15] Our model assumes that the distribution of the log unobserved auction-specific component entering bank valuations has mean zero, $\mathbb{E}[\ln Y] = 0$, thereby fixing this location term and yielding identification of the distributions of X and Y.

5 Estimation

Estimating the structural parameters follows the steps described in identification. First we estimate $G(\cdot)$, the distribution of bank response times to being pushed OUT with a kernel estimator using data on when a bid was pushed OUT and the next time the affect bank bid, either OUT, ONM or INM. Then we discretize the likelihood (4) to estimate $F_V(\cdot)$, the distribution of the valuations, with data on the ONM rates and whether or not banks submitted INM bids when they had the opportunity. The joint distribution $F_{V_1,V_2}(\cdot)$ is also estimated this way with observations formed on pairs of banks bidding in the same auction. Finally, deconvoluting the estimates of $\phi_{\ln X}(t)$ and $\phi_{\ln Y}(t)$ with an inverse Fourier transform yields estimates for the distribution functions $F_X(\cdot)$ and $F_Y(\cdot)$. This section elaborates each step, Appendix D.1 establishes the consistency of the identified portions of $F_V(\cdot)$ and by extension $F_{V_1,V_2}(\cdot)$, Appendix D.2 describes the implementation of the estimation in further detail, while Appendix D.3 explains the subsampling procedure used to obtain the estimated asymptotic errors.

[15] One of the assumptions needed to apply the result of Kotlarski (1966) is that the characteristic functions are non-vanishing. This follows from our assumption of bounded support for the valuation distribution $\mathcal{V}$; Krasnokutskaya (2011) provides a proof in a similar context.

5.1 Bid Monitoring Opportunity Arrival Rate $G_t(\cdot)$

The probability distribution function governing the arrival of bid opportunities, $G_t(\cdot)$, is estimated off the observations of bank reaction times to being pushed OUT. We used a kernel smoothing function to estimate this nonparametric regression function because time $t \in \{1, \ldots, 1800\}$, measured in seconds, is ordinal and has a fine grid. Recalling ρ_{nj} is when bank n receives its j^{th} opportunity to bid and $\tau_{nj} > \rho_{nj}$ is when its j^{th} bid falls OUT (where $\tau_{n0} \equiv 0$), we estimate $G_t(z)$pointwise for each (t, z) with

$$\widehat{G}_t(z) = \frac{\sum_{n=1}^{N} \sum_{j=0}^{J_n - 1} K\left(\frac{\tau_{nj} - t}{h}\right) \mathbb{1}\{\rho_{n,j+1} - \tau_{nj} < z\}}{\sum_{n=1}^{N} \sum_{j=1}^{J_n} K\left(\frac{\tau_{nj} - t}{h}\right)}$$

where $K(\cdot)$ is a normal kernel and the bandwidth parameter h is set by Silverman's rule-of-thumb. Differentiating $\widehat{G}_t(z)$ with respect to z yields an estimate of its density $\widehat{g}_t(z)$.

5.2 Distribution of Valuations $F_V(\cdot)$ and $F_{V_1,V_2}(\cdot)$

The estimation of $F_V(\cdot)$ is based on substituting estimates of $\widehat{G}_t(\cdot)$ into the portion of the likelihood equation (4) that depends on $F_V(\cdot)$ and maximizing the resulting criterion function on a grid of 30 points, which we now denote by $\mathbf{V} \subseteq \mathcal{V}$, that includes the points $\underline{v} = 0$ and $\overline{v} = 1$. The criterion function is the product over each bank's individual contribution:

$$[1 - F_V(r_{\rho_J})] + \mathbb{1}\{\widetilde{W}_{b_J}\} H(F_V, \hat{G}_{\tau_J}) \tag{7}$$

The joint distribution $F_{V_1,V_2}(\cdot)$ is estimated using an analgous procedure, by substituting our estimator for $\widehat{G}_t(z)$ into $(C.4)$ to form the criterion function.

5.3 Distribution of Multiplicative Factors in Valuations $F_X(\cdot)$ and $F_Y(\cdot)$

The estimator for $\phi_{\ln V_1,\ln V_2}(t_1,t_2)$, $\phi_{\ln V_1,\ln V_2}(t_1,t_2)$, the characteristic function for $(\ln V_1, \ln V_2)$, is given by its discrete analogue, computed directly off $\hat{F}_{V_1,V_2}(\cdot)$ as:

$$\hat{\phi}_{\ln V_1,\ln V_2}(t_1,t_2) = \sum_{v_1\in\mathbf{V}}\sum_{v_2\in\mathbf{V}} e^{(it_1 \ln v_1 + it_2 \ln v_2)}\hat{p}_{\ln V_1,\ln V_2}(\ln v_1, \ln v_2)$$

where $\hat{p}_{\ln V_1,\ln V_2}(\ln v_1, \ln v_2)$ is the estimated joint probability of drawing banks with valuations v_1 and v_2 from the same auction. Noting that we have normalized $\mathbb{E}[\ln Y] = 0$, estimates for the characteristics functions $\phi_{\ln X}(t)$ and $\phi_{\ln Y}(t)$ are generated from $\phi_{\ln V_1,\ln V_2}(t_1,t_2)$ using (5) and (6):

$$\begin{aligned}
\hat{\phi}_{\ln Y}(t) &= \exp\left(\int_0^t \frac{\partial}{\partial u}\left[\frac{\hat{\phi}_{\ln V_1,\ln V_2}(u,v)}{\hat{\phi}_{\ln V_1,\ln V_2}(u,0)\hat{\phi}_{\ln V_1,\ln V_2}(0,v)}\right]_{u=0} dv\right) \\
\hat{\phi}_{\ln X}(t) &= \frac{\hat{\phi}_{\ln V_1,\ln V_2}(t,0)}{\hat{\phi}_{\ln Y}(t)}
\end{aligned}$$

The last step recovers estimates for $F_X(\cdot)$ and $F_Y(\cdot)$ from the characteristic functions $\hat{\phi}_{\ln X}(t)$ and $\hat{\phi}_{\ln Y}(t)$. Using the inverse Fourier transform for discrete random variables we obtain the estimates for $\hat{p}_{\ln Y}$ and $\hat{p}_{\ln X}$, the probability mass functions for X and Y:

$$\begin{aligned}
\hat{p}_{\ln Y}(kh_Y) &= \frac{1}{2\pi/h_Y}\int_{-\pi/h_Y}^{\pi/h_Y} e^{itkh_Y}\hat{\phi}_{\ln Y}(t)dt \\
\hat{p}_{\ln X}(kh_X) &= \frac{1}{2\pi/h_X}\int_{-\pi/h_X}^{\pi/h_X} e^{itkh_X}\hat{\phi}_{\ln X}(t)dt
\end{aligned}$$

where k takes integer values and the random variable takes values kh for some $h > 0$. Finally, a monotone transformation of the estimated distributions of $\ln Y$ and $\ln X$ yields estimates of $F_Y(\cdot)$ and $F_X(\cdot)$.

6 Results

Figure 4 shows the estimated bid arrival distributions for before and after the 2008 crisis. As Figure 3b foreshadowed, banks respond more quickly to being pushed OUT later in the auction than earlier. Practically all the bids pushed OUT at 25 minute mark were reviewed within 250 seconds, but in the post-2008 period more than 10 percent of those pushed out at the 10 minute mark had not been reviewed within that time frame. The key difference between the circled plot in Figure 3b and plots of both distributions in the bottom panel of Figure 4 is that the former plateaus at about 0.65 (compared to almost 1.00). This difference is explained by the fact that upon discovering their previous bid is OUT, a high percentage of banks withdraw from the auction rather than bidding INM. Although the bid response times to being pushed OUT are almost indistinguishable towards the end of the auction, both statistically and in magnitude, there are noticeable differences at the beginning of the auction: response times are faster in the pre-2008 auctions.

Figure 5 shows the estimated distributions of private values $F_X(\cdot)$ and auction-level unobserved heterogeneity $F_Y(\cdot)$ for the pre-and-post-2008 periods.[16] The top figure shows the distribution of the private valuation component X for the pre-2008 period has a higher mean than the post (0.529 versus 0.274). More generally, the point estimates suggest the distribution of the private valuations in pre-2008 first order stochastically dominate the post-2008 distribution.

The role of unobserved heterogeneity between auctions is also greater in post-2008 period than in the earlier period: noting that the mean of the log of this random variable is zero by construction, the bottom figure exhibits a mean preserving compression of the distribution

[16]Recall that the overall valuation, expressed as an interest rate, takes the form $\tilde{v} = x * y + \underline{r}$ where x is the private valuation, y is the auction-specific factor common to all bidders within the auction, and $\underline{r}$ is the auction reserve interest rate. Figure 5 (a) shows the distribution of the private valuation and Figure 5 (b) shows the distribution of the auction-specific component.

when shifting from the pre-2008 period to the post-2008 period. However we qualify these remarks by noting that the evidence on this point is weak: the estimated probability density function for each period, both pre-2008 and post-2008, fits within the 95 percent confidence bounds of the other period.

The estimated differences between the distribution functions plotted in Figures 4 and 5 inferred from the model can be plausibly reconciled with the facts presented in Table 1. When the distribution of valuations shifted to the left after the crash (as depicted in Figure 5), demand for CD's fell, fewer banks bid in the auctions, there were more unfilled auctions (in which every bidding bank won), and a higher proportion of bidding banks won (documented in Rows I, XVI and VII of Table 1). Three factors combine to reduce monitoring incentives post-2008. First, banks are less likely to draw valuations worth bidding on. Second, they face reduced competition in the (less likely) event of drawing a high valuations. Finally, they learn less about the distribution of rival bids by repeatedly bidding just above the ONM rate, as opposed to making an early jump bid. This explains why Figure 4 shows response times that are slower in the early phases of auctions conducted post-2008. Reduced monitoring led in turn to placing (higher) INM bids less likely to fall OUT later in the auction, thus reducing the total number of bids placed per parcel within an auction, and raising the prospect of money left on the table, both predictions borne out in Rows II and XIV.

Figure 4: Bid Opportunity Arrival Time Distributions $G_t(z)$

(a) 10 minutes into auction

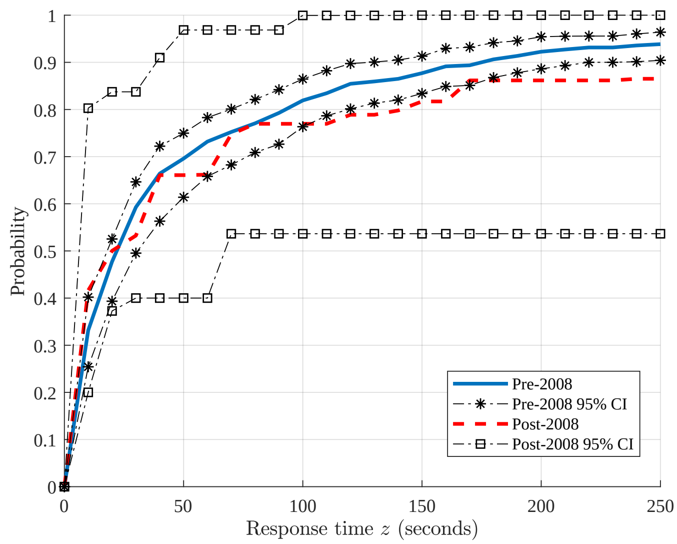

(b) 25 minutes into auction

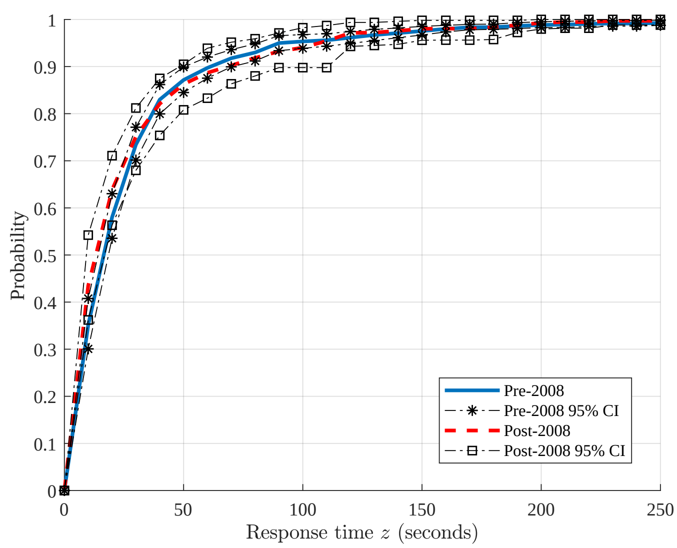

Note: This figure shows the estimated distribution of time to receive another bidding opportunity after being pushed OUT at two different points in the auction. Figure (a) shows the distribution for a bank pushed OUT at ten minutes into the auction and Figure (b) shows the distribution for a bank pushed out at 25 minutes into the auction. Response times to being pushed OUT at ten minutes can be several minutes or longer, with slightly higher average response times for banks in the post-2008 period. At later stages in the auction banks in both the pre-2008 and post-2008 samples respond quickly to being pushed OUT.

Figure 5: Distribution of Private Values and Unobserved Heterogeneity

(a) Private values

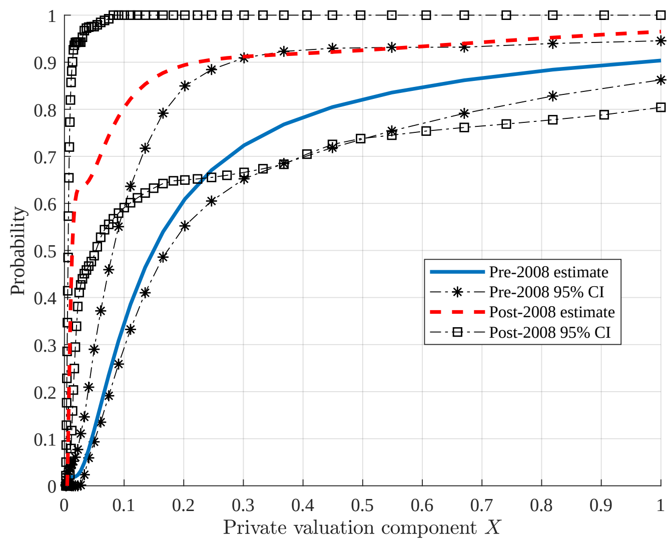

(b) Unobserved heterogeneity

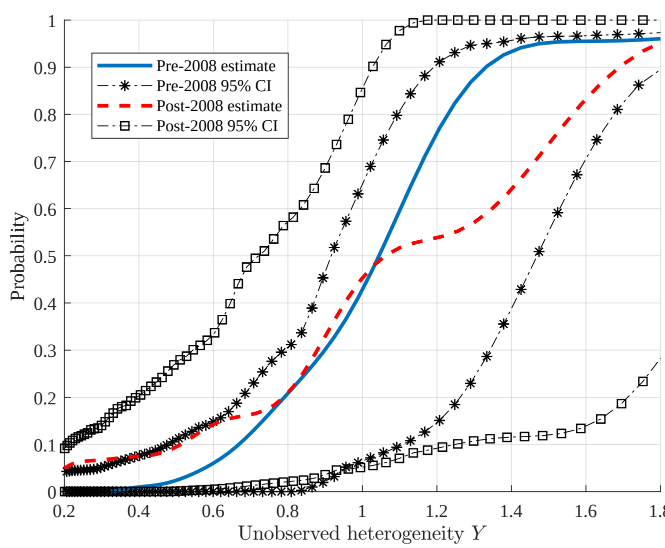

Note: Figure (a) shows the distributions of banks' private valuations X and Figure (b) shows the distribution of auction-specific unobserved heterogeneity Y. Each bank's overall valuation is an interest rate expressed as $\tilde{V} = XY + \underline{r}$ where $\underline{r}$ is the auction reserve rate. The term XY represents the maximum percentage points above the reserve rate that a bank is willing to pay for funds.

7 Costs of Bidding Frictions

Bidding frictions are costly because banks cannot always react to rival bids before the end of the auction. Consequently the set of winning bids does not necessarily correspond to the banks with the highest valuations, because some banks with high valuations may never have the chance to respond to competitive bids. This section analyzes the potential for reallocation and quantify revenue losses to sellers attributable to frictions and/or opportunity cost.

In principle the set of equilibria could be computed using the structural estimates but there are two practical barriers to this approach. The first is computational: the number of equilibria is not determined and each equilibrium is a complicated non-stationary function in continuous time. Secondly the mixture of equilibria played by banks is not known. Instead, we derive bounds on the expected value of winning banks that apply regardless of which equilibrium is selected by banks.

The bounds are based on two simple ideas: first, that a bank's valuation is higher than all of its bids, and second that if a bid falls OUT the bank would raise its bid to at least the current ONM rate at the next bidding opportunity if its valuation was high enough. Let W denote the event of placing a winning bid and $\tilde{W}$ its complement, losing. Appealing to the law of iterated expectations, $\mathbb{E}[V]$ can be expressed as a weighted sum of $\mathbb{E}[V|W]$ and $\mathbb{E}[V|\tilde{W}]$. Upon rearrangement we obtain:

$$\mathbb{E}[V|W] = \frac{\mathbb{E}[V] - \Pr(\tilde{W})\mathbb{E}[V|\tilde{W}]}{\Pr(W)}. \tag{8}$$

We bound $\mathbb{E}[V]$ as follows. Following the notation of Section 4, denote by $\{t_s\}_{s=1}^{S}$ the times at which the ONM rate increases, set $t_{S+1} \equiv T$, and let $\{r_s\}_{s=1}^{S}$ denote the ONM rates at those times. Also let r_{τ_J} denote the the ONM rate that pushes bank's final bid OUT at time $\tau_J \equiv t_{S_J}$, for some index value $S_J \in \{1, \ldots, S\}$. Since the bank would bid the first opportunity after its previous bid falls OUT if its valuation remains higher than the ONM rate:

$$\mathbb{E}[V|\tilde{W}] < \left(\sum_{s=S_J}^{S} \frac{G_{\tau_J}(t_s) - G_{\tau_J}(t_{s+1})}{G_{\tau_J}(T)} \int_{r_{\tau_J}}^{r_s} \frac{v f_V(v)}{F_V(r_s) - F_V(r_{\tau_J}))} dv \right) \tag{9}$$

where $f_V(\cdot)$ is the probability density function (or discrete probabilities within the relevant interval) of V. To derive this inequality, note the integrals from r_{τ_J} to r_s with respect to v are conditional expectations of V when V lies between those two ONM rates; they are weighted by the probability that the bank has an opportunity to bid between t_s and t_{s+1} relative having any opportunity between t_J and T. To bound $\mathbb{E}[V|\tilde{W}]$ from below, define $\tau_{J-1} \equiv t_{S_J-1}$, which implies $r_{\tau_{J-1}}$ the ONM rate immediately preceding r_{τ_J}. Since $v \geq r_{\tau_{J-1}}$it now follows that:

$$\mathbb{E}[V|\tilde{W}] > \left(\sum_{s=S_J-1}^{S} \frac{G_{\tau_{J-1}}(t_s) - G_{\tau_{J-1}}(t_{s+1})}{G_{\tau_{J-1}}(T)} \int_{r_{\tau_{J-1}}}^{r_s} \frac{v f_V(v)}{F_V(r_s) - F_V(r_{\tau_{J-1}}))} dv \right) \tag{10}$$

The quantiles of $F(\cdot)$ are identified at the ONM rates, yielding upper and lower bounds on $\mathbb{E}[V|W]$ as well as the conditional expectations terms in (9) and (10). Also $\Pr(W)$ is identified from the proportion of banks that win, as is $G_t(\cdot)$. In this way we bound (8) on both sides. Substituting these expressions into (8) yields the upper and lower bounds on $\mathbb{E}[V|W]$.

Table 2: Efficiency Measurements

	Pre-2008	**Post-2008**
Lower Bound on $\mathbb{E}[V|W]$	0.57	0.33
Upper Bound on $\mathbb{E}[V|W]$	0.59	0.36
Expected Valuation of Winner, Uniform Price	0.63	0.36
% Increase in Revenue using Uniform Price	8.54	2.63

Table 2 lists the expected valuation of winning banks for the pre-2008 and post-2008 auctions and compares these values with the expected valuations of winners from a frictionless uniform price auction mechanism in which the price paid by all winning banks is the highest loser's valuation. The bounds are computed assuming the set of ONM rates cover the entire set of valuations, essentially because the ONM rates saturate the real line in some practical sense. For example over 99 percent of all bids are at or below the highest ONM rate observed in the data. The bounds are quite tight. The efficiency gains for the pre-2008 auctions are between 0.04 and 0.06 percentage points, while the post-2008 gains are between 0 and 0.03 percentage points. Hence the costs of frictions are higher in the pre-2008 auctions. This difference has two main causes: there is greater variation in pre-2008 bank valuations, and so the probability of a very high valuation bank being left OUT due to frictions is higher prior to the financial crisis, and (from Table 1) fewer banks compete post-2008, so fewer are pushed OUT, and a greater proportion of the auctions are unfilled.

The table also compares the revenue generated from the frictionless uniform price mechanism with the actual revenue generated by the auction platform. Eliminating frictions and using uniform pricing would have increased auction revenue by 8.54 percent prior to the financial crisis and 2.63 percent after. Auction revenue is limited because high value banks sometimes cannot update their bids after being pushed OUT. The revenue losses are smaller after the financial crisis for the same reason as the efficiency increases: compressed valuations and more unfilled auctions post-2008.

The choice of auction mechanism is puzzling given the estimated losses due to bidding frictions. Why should the auctioneer not run a sealed-bid uniform price auction to allocate funds? We suggest two reasons why the current mechanism may be preferred to such an alternative mechanism despite the losses due to frictions: collusion and corruption. Recent theoretical work has demonstrated the potential for collusive outcomes in similar games with more public information on player actions. Kamada and Kandori (2017) study a continuous-time game with a deadline and random arrival of opportunities to change one's action. In their model, all actions are perfectly observable, and players are able to coordinate to obtain results very near the full-collusion outcome. The coordinated equilibria of the type demonstrated in Kamada and Kandori (2017) may be one reason why the auction platform opted for providing banks with limited information. Additionally, the ascending-price mechanism may alleviate concerns of corruption. If the mechanism were sealed bid, it would be possible for the auctioneer to collect information on banks' bids and give it to one of their competitors in exchange for a side payment. This type of interaction is more difficult in an ascending auction where banks can update their bids thus reducing the value of this type of confidential information.

8 Conclusion

In this paper we study ascending price auctions for financial products in local markets and provide evidence that banks in these auctions face frictions. To analyze the market we build a model of bidding in ascending price auctions where bidding opportunities are not automatic or at will, but arrive stochastically. These opportunities reflect alternative uses of time in managing a broad portfolio of assets. The model may have multiple equilibria, and the set of equilibrium strategies may contain strategies that are non-monotone in bank valuations, such as pooling strategies, and mixed strategies.

We do not restrict attention to a single pure strategy equilibrium. Instead, we pursue an approach that is robust to different equilibria being played across the auctions in the data, achieving identification through a restriction to undominated strategies. Through this restriction and the identification of the bidding opportunity arrival time distribution, we are able to identify both the distribution of private values as well as the distribution of unobserved auction-level heterogeneity up to quantiles defined by all the ONM bids. We also separate the valuation distribution into a private component independently distributed across banks, and an auction specific factor that is common to all the banks in the same auction. These two component distributions are estimated off the joint distribution of pairs of valuations within an auction; our approach differs from previous work that sought to point-identify individual bank valuations to estimate the characteristic function of the joint distribution directly.

The structural estimates are used to quantify the cost of bidding frictions and to assess how the 2008 financial crisis impacted banks' demands for deposit funds. We find that bidding frictions are costly from both a welfare and revenue standpoint. The expected valuation of winning banks could be improved by as much as 0.06 percentage points in pre-2008 auctions and 0.03 percentage points in post-2008 auctions. Auction revenue is 8.54 percent and 2.63 percent higher in a uniform price auction with no bidding frictions in pre-2008 and post-2008 auctions, respectively. Our results underscore practical measures taken to alleviate bidding frictions in many auction platforms, such as extending the auction clock until bidding activity ceases.

Acknowledgments: We thank Timothy Derdenger for pointing us to this data, and John Carver and Jon Smith at Grant Street for providing it. We thank the guest editors and two anonymous referees for helpful feedback and guidance. For detailed comments, we thank Dan Ackerberg, Lanier Benkard, Phil Haile, Ken Hendricks, Ali Hortacsu, Isabelle Perrigne, Martin Pesendorfer, John Rust, Quang Vuong, and Ariel Pakes. We have also benefitted from presentations at Econometric Society Meetings in Asia, the Implementation of Structural Dynamic Models conference at UCL, and the CRES Applied Microeconomics Conference at WashU, the Asia Pacific IO conference, and at economics seminars in UT Austin, LSE, Princeton, Johns Hopkins, Vanderbilt, CEMFI, Carlos III, Autonoma, UCLA, Yale, Georgetown, Rice, Duke, Penn State, and UC Berkeley.

References

Ackerberg, D., K. Hirano, and Q. Shahriar (2016): "Identification of Time and Risk Preferences in Buy Price Auctions," *Quantitative Economics*, Conditionally Accepted.

Ambrus, A., Y. Ishii, and J. Burns (2013): "Gradual Bidding in eBay-like Auctions," Discussion paper.

Aradillas-López, A., A. Gandhi, and D. Quint (2013): "Identification and Inference in Ascending Auctions with Correlated Private Values," *Econometrica*, 81(2), 489–534.

Aradillas-Lopez, A., and E. Tamer (2008): "The identification power of equilibrium in simple games," *Journal of Business & Economic Statistics*, 26(3), 261–283.

Athey, S., and P. Haile (2007): "Nonparametric Approaches to Auctions," *Handbook of Econometrics*, 6(Part A), 3847–3965.

Athey, S., and P. A. Haile (2002): "Identification of Standard Auction Models," *Econometrica*, 70(6), 2107–2140.

Athey, S., J. Levin, and E. Seira (2011): "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions," *The Quarterly Journal of Economics*, 126(1), 207–257.

Avery, C. (1998): "Strategic Jump Bidding in English Auctions," *The Review of Economic Studies*, 65(2), 185–210.

Bajari, P., and A. Hortacsu (2003): "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, pp. 329–355.

Cassola, N., A. Hortaçsu, and J. Kastl (2013): "The 2007 Subprime Market Crisis Through the Lens of European Central Bank Auctions for Short-Term Funds," *Econometrica*, 81(4), 1309–1345.

Decarolis, F. (2017): "Comparing Procurement Auctions," *International Economic Review*, Forthcoming.

Easley, R. F., and R. Tenorio (2004): "Jump Bidding Strategies in Internet Auctions," *Management Science*, 50(10), 1407–1419.

Freyberger, J., and B. J. Larsen (2017): "Identification in Ascending Auctions, with an Application to Digital Rights Management," Discussion paper, National Bureau of Economic Research.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68(3), 525–574.

Haile, P. A., and E. Tamer (2003): "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, 111(1), 1–51.

Hollifield, B., R. A. Miller, and P. Sandås (2004): "Empirical analysis of limit order markets," *The Review of Economic Studies*, 71(4), 1027–1063.

Hollifield, B., R. A. Miller, P. Sandås, and J. Slive (2006): "Estimating the Gains from Trade in Limit-Order Markets," *The Journal of Finance*, 61(6), 2753–2804.

Hong, H., and M. Shum (2003): "Econometric Models of Asymmetric Ascending Auctions," *Journal of Econometrics*, 112(2), 327–358.

Hörner, J., and N. Sahuguet (2007): "Costly Signalling in Auctions," *The Review of Economic Studies*, 74(1), 173–206.

Kamada, Y., and M. Kandori (2017): "Revision Games," Discussion paper.

Kline, B. (2018): "An empirical model of non-equilibrium behavior in games," *Quantitative Economics*, 9(1), 141–181.

Kline, B., and E. Tamer (2012): "Bounds for best response functions in binary games," *Journal of Econometrics*, 166(1), 92–105.

KOTLARSKI, I. (1966): “On some characterization of probability distributions in Hilbert spaces,” *Annali di Matematica Pura ed Applicata*, 74(1), 129–134.

KRASNOKUTSKAYA, E. (2011): “Identification and Estimation of Auction Models with Unobserved Heterogeneity,” *The Review of Economic Studies*, 78(1), 293–327.

LI, Q. (1996): “Nonparametric Testing of Closeness Between Two Unknown Distribution Functions,” *Econometric Reviews*, 15(3), 261–274.

LI, T., AND Q. VUONG (1998): “Nonparametric Estimation of the Measurement Error Model using Multiple Indicators,” *Journal of Multivariate Analysis*, 65(2), 139–165.

MILGROM, P. R., AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, pp. 1089–1122.

PAARSCH, H. J. (1997): “Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales,” *Journal of Econometrics*, 78(1), 333–357.

POLITIS, D. N., J. P. ROMANO, AND M. WOLF (1999): *Subsampling*. Springer, New York.

RAO, B. P. (1992): *Identifiability in Stochastic Models: Characterization of Probability Distributions*. Academic Press.

ZEITHAMMER, R., AND C. ADAMS (2010): “The Sealed-bid Abstraction in Online Auctions,” *Marketing Science*, 29(6), 964–987.

A Test of Bidding Frictions

To test whether bidding frictions exist, let $f_{\underline{W}}(b)$ denote the probability density function of the lowest winning bid, and $f_{\overline{W}}(b)$ the density of highest winning bid. Under the null hypothesis of no frictions the distributions of the lowest winning bid and highest winning bid are identical: $H_0 : f_{\underline{W}}(\cdot) = f_{\overline{W}}(\cdot)$. Let $b_i^{\underline{W}}$ denote the lowest winning bid from the i^{th} auction and $b_j^{\overline{W}}$ the highest winning bid from the j^{th} auction. Assume all observations are independent. Following Li (1996) we base the test statistic for H_0 on $\int(\hat{f}_{\underline{W}}(b) - \hat{f}_{\overline{W}}(b))^2 db$. Define:

$$T_{diff} \equiv Nh^{1/2} I_N / \sqrt{\hat{\sigma}^2}$$

where N is the sample size for both winning bids and losing bids, h is the bandwidth parameter set by Silverman's rule of thumb, $K(\cdot)$ is the standard normal density function, and:

$$I_N = \frac{1}{hN(N-1)} \left[\sum_{i=1}^{N} \sum_{\substack{j=1,\\ j\neq i}}^{N} K(\frac{b_i^{\underline{W}} - b_j^{\underline{W}}}{h}) + \sum_{i=1}^{N} \sum_{\substack{j=1,\\ j\neq i}}^{N} K(\frac{b_i^{\overline{W}} - b_j^{\overline{W}}}{h}) - 2 \sum_{i=1}^{N} \sum_{\substack{j=1,\\ j\neq i}}^{N} K(\frac{b_i^{\underline{W}} - b_j^{\overline{W}}}{h}) \right]$$

$$\hat{\sigma}^2 = \frac{2}{N^2} \sum_{i=1}^{N} \sum_{j=1,}^{N} \left[K(\frac{b_i^{\underline{W}} - b_j^{\underline{W}}}{h}) + K(\frac{b_i^{\overline{W}} - b_j^{\overline{W}}}{h}) + 2K(\frac{b_i^{\underline{W}} - b_j^{\overline{W}}}{h}) \right] \int K^2(u) du$$

Li (1996) proves $T_{diff} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1)$ under H_0. Since the value of the test statistic is 13.35, we reject the null hypothesis of equality between the lowest and highest winning bid distributions at the 1 percent confidence level.[17]

[17] Zeithammer and Adams (2010) provide similar tests of the "sealed-bid abstraction" in online markets and also reject these models.

B Independence of Bidding Opportunities and Valuations

Our structural model assumes bidding opportunities to be independent of bank valuations. To investigate how plausible this assumption is, this section presents two pieces of descriptive evidence on observed bidding patterns. For reasons explained below, the maximum bid submitted by losing banks is a reasonable proxy for their valuation. Appendix B.1 documents the relationship between those bids and reaction times to being pushed OUT, by examining the non-parametrically estimated reaction times of banks to being pushed OUT with two minutes remaining in the auction for different maximum bid levels. Then in Appendix B.2 we present regression results from regressing the reaction time against the maximal bid, for all banks and also for the subset of banks whose final bid is OUT, accounting for other factors that affect reaction times.

Both diagnostics provide scant evidence for the view that banks adjust their monitoring rates to their valuation. We conclude that it is more likely the monitoring technology is set at the institution level, for example through long-standing staffing arrangements, and not quickly adjusted to increase or decrease monitoring intensity for one particular auction or valuation draw.

B.1 Empirical distribution of reaction times

Section 2 shows that banks making INM bids first tend to creep up to the ONM rate and only then jump. From an institutional perspective, this is a sensible bidding strategy because it reveals the ONM rate to the bank, one of the few indicators it can use to gauge the intensity of the competition, and economizes on the size of the bid increment necessary to restore being INM. Following this strategy leads banks to stop bidding when they reach their valuation if it is less than the current ONM rate. Only in this case is their final bid OUT. For these reasons we believe a bank's final bid is a useful proxy for its valuation if that bid is OUT.

Figure 6 shows the nonparametrically estimated cumulative distribution function of the time elapsed to its next bid when a bank is pushed OUT two minutes before the auction ends as bid activity enters the most intense phase, conditional on the bank's final bid being OUT. Thus the horizontal axis measures the time OUT preceding the bank's next bid when two minutes remain before the auction ends. The vertical axis shows the percentile of banks that have responded by that time. We order the banks by their final (maximum OUT) bid, forming the 90th, 50th, and 10th percentiles in that distribution. Figure 6 plots the estimated distribution of elapse times for these three quantiles.

The reaction times of banks whose highest bid is in the 90th percentile hardly differ from those in the 10th percentile. Indeed, the distribution functions cross each other multiple times. Furthermore, all three distributions (10th, 50th, and 90th percentiles of OUT bids) are similar to the unconditional distribution of reaction times at two minutes for all banks.

B.2 Reaction time regressions

Table 3 reports results from regressing the time elapsed until the next bid on the maximum bid within an auction, controlling for different times in the auction, size of the bid parcel, and the onset of the financial crisis (indicated by the time dummy for post 2008). The dependent variable is the time elapsed in seconds between a bid being pushed out of the money and the bank responding by placing a new bid. The maximum bid of every bank in every auction is included in the first three specifications. In the other three specifications only those banks whose final bid was OUT are included.

The effect of higher maximum bids is insignificant across all specifications. Moreover the negative point estimates have a small quantitative impact. For example the point estimate from column (2), implies that increasing the maximum bid by a standard deviation increase

Figure 6: Time Elapsed to Next Bid, Two Minutes Remaining for OUT Banks

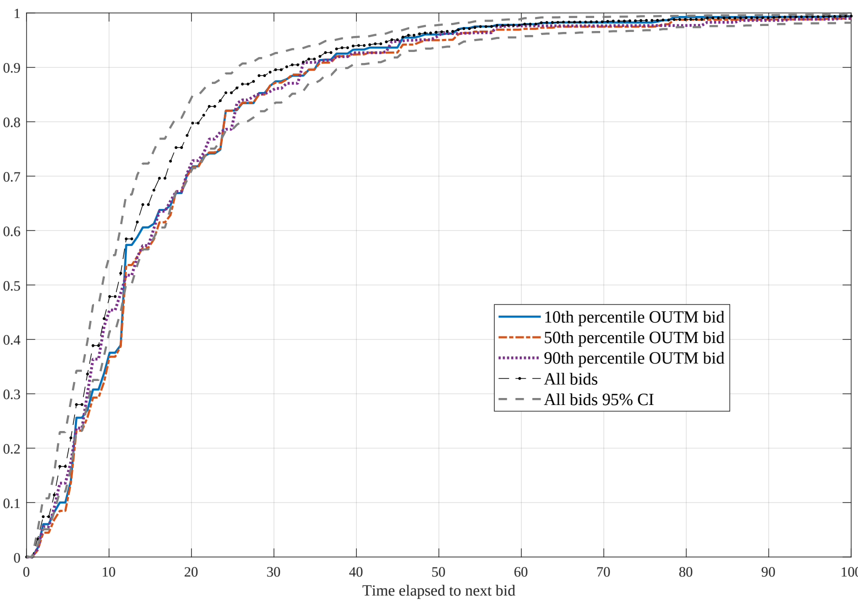

reduces the response time by approximately one second. Finally, focusing on the regressions in columns (4) through (6) where we are most confident that the maximal bids are closely related to the valuations, the sign, albeit insignificant, is counterintuitive: longer response times are associated with higher losing creeping bids.

C Identification and the Likelihood

Appendix C.1 presents the general form of the likelihood in which the hazard rate governing bidding arrival times may depend on valuations, simplifies it to the case used in identification and estimation where the hazard rate is independent of valuations, and discusses the implications of relaxing the assumption for identification and estimation. Appendix C.2 demonstrates that for any finite partition of the valuation space the likelihood function is globally concave,

Table 3: Regression results for maximum bid on bank response time

	All bidders			Out-of-the-money bidders		
	(1)	(2)	(3)	(4)	(5)	(6)
Maximum Bid	-3.293	-3.483	-4.447	10.36	10.57	12.16
	(-0.92)	(-0.98)	(-1.26)	(1.52)	(1.56)	(1.79)
Time pushed out	-0.0490***	-0.0484***	-0.0475***	-0.0458***	-0.0465***	-0.0471***
	(-16.43)	(-15.71)	(-15.37)	(-8.28)	(-7.99)	(-8.09)
$\mathbb{1}\{\text{Post - 2008}\}$		-2.311	-5.137		3.628	4.110
		(-0.83)	(-1.82)		(0.66)	(0.74)
Quantity			0.00209***			-0.00116
			(5.90)			(-1.93)
Constant	71.03***	70.86***	62.71***	77.15***	77.34***	80.45***
	(13.83)	(13.73)	(11.49)	(7.61)	(7.58)	(7.75)
N	11970	11970	11970	3701	3701	3701
R^2	0.027	0.027	0.030	0.032	0.033	0.033

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The variable "Max bid" is the highest bid submitted by each individual bank within that auction, and "Time pushed out" represents that time at which the bank's previous bid was pushed out of the money by another bank. The "Max bid" variable is postively associated with bank valuations, as higher maximum bids within an auction should be correlated with higher valuations.

yielding identification of the cumulative distribution function of valuations over the quantiles at each observed ONM rate. Then in Appendix C.3 we show that these quantiles converge to the true distribution as the partition becomes increasingly fine.

C.1 The likelihood for bidding opportunities and valuations

Following the main text, assume the valuation of each bank and its bid arrival times are independent. Then the joint likelihood for the arrival of bidding opportunities and the distribution of a bank's decisions to continue bidding or not can be expressed as:

$$\begin{aligned}
\mathcal{L}(F_V, G_\tau) &= \prod_{j=1}^{J} g_{\tau_{j-1}}(\rho_j)\left[\frac{1-F_V(r_{\rho_j})}{1-F_V(r_{\rho_{j-1}})}\right]\times \\
&\left(1+\mathbb{1}\{\widetilde{W}_{b_J}\}\left\{\sum_{s=S_J}^{S}[G_{\tau_J}(t_{s+1})-G_{\tau_J}(t_s)]\frac{F_V(r_{t_s})-F_V(r_{\rho_J})}{1-F_V(r_{\rho_J})}+\frac{(1-G_{\tau_J}(T))}{1-F_V(r_{\rho_J})}\right\}\right) \\
&= \prod_{j=1}^{J} g_{\tau_{j-1}}(\rho_j)\times\left(\begin{array}{l} [1-F_V(r_{\rho_J})]+\mathbb{1}\{\widetilde{W}_{b_J}\}\{(1-G_{\tau_J}(T))\} \\ +\mathbb{1}\{\widetilde{W}_{b_J}\}\sum_{s=S_J}^{S}[G_{\tau_J}(t_{s+1})-G_{\tau_J}(t_s)]\,[F_V(r_{t_s})-F_V(r_{\rho_J})]\end{array}\right)
\end{aligned} \tag{C.1}$$

The product over j in the first equality is formed from (3), while the intuition for the expression following $\widetilde{W}_{b_J}$ is discussed where $H(F_V, G_{\tau_J})$ is defined in (2). The second equality appeals to the normalization that $F_V(r_{\rho_0}) = 0$, and telescopes the product terms involving the value quantiles. Thus the numerator from the j^{th} term is canceled by the denominator from term $j+1$. The numerator from the J^{th} and final term is then distributed into the remaining two terms inside the sum. Substituting (2) into the summation term gives (4) in the text.

If the independence assumption for valuations and bid arrival times is violated, then the top line of $\mathcal{L}(F_V, G_\tau)$ generalizes to:

$$\begin{aligned}
&\prod_{j=1}^{J}\int_{r_{\rho_j}}^{\overline{v}} g_{\tau_{j-1}}(\rho_j\,|v)\frac{f_V(v)}{1-F_V(r_{\rho_{j-1}})}dv \\
&\times\left(1+\frac{\mathbb{1}\{\widetilde{W}_{b_J}\}}{1-F_V(r_{\rho_J})}\left\{1+\sum_{s=S_J}^{S}\int_{r_{\rho_J}}^{r_{t_s}}[G_{\tau_J}(t_{s+1}\,|v)-G_{\tau_J}(t_s\,|v)]f_V(v)dv-\int_{r_{\rho_J}}^{\overline{v}} G_{\tau_J}(T\,|v)f_V(v)\,dv\right\}\right)
\end{aligned} \tag{C.2}$$

Comparing (C.2) with (C.1), there are two complications associated with relaxing the independence assumption. Both arise from a feature of (C.1) that (C.2) lacks: a limited information likelihood for $G_\tau(\cdot)$ that only makes use of the time intervals $\rho_j - \tau_{j-1}$, between falling OUT and having another opportunity to bid. The first is in establishing identification. The presence of conditioning terms in the distribution of bid arrival times in $g_{\tau_{j-1}}(\rho_j \,|v\,)$ negates the decomposition of (C.2) into two pieces that lead to the identification of $G_t(\cdot)$ from the monitoring time intervals alone; this is compounded by the fact that $F_V(\cdot)$ can be identified from last bids and the ONM rates that follow them if $G_\tau(\cdot)$ is separately identified. Therefore a natural identification strategy under the independence assumption fails when it is relaxed, leaving no clear alternative approach. The second is computational. Under (C.1) but not under (C.2) sequential estimation is possible.

In principle both problems might be overcome by imposing parametric restrictions on $G_t(\cdot)$ and $F_V(\cdot)$. However Appendix B shows there is little evidence to suggest bank valuations have a substantial impact on bid arrival time distributions. For these reasons we opted for a nonparametric approach that imposed the independence assumption. Falsely imposing the independence assumption might lead to underestimating the proportion of banks with lower valuations. These banks tend to lose more often than banks with high valuations, so we might expect their monitoring rates to be lower. The independence assumption attributes bank inactivity to the bad luck of not having another bidding opportunity. If in reality low valuation banks have fewer bidding opportunities, our identification strategy undercounts the relative frequency with which they withdraw from the auction when compelled to choose.

C.2 Likelihood for pairs of banks

In order to decompose the probability distribution of bank valuations $F_V(\cdot)$ into its idiosyncratic $F_X(\cdot)$ and auction $F_Y(\cdot)$ components, we first identify and estimate define the joint distribution for any two valuations of any banks competing in the same auction, which we now denote by $F_{V_1,V_2}(\cdot)$. The arguments above establish that under the maintained assumption of valuations and monitoring times $G_t(\cdot)$ is identified from the time intervals $\rho_j - \tau_{j-1}$. Moreover as we establish below, $G_t(\cdot)$ can be nonparametrically estimated using limited information maximum likelihood methods. We exploit these results by showing that $F_{V_1,V_2}(\cdot)$ can be identified and estimated from data on the status of their final bid, whether it is a winning bid or not.

Adapting the notation established for analyzing each bank separately, let J_i denote the last time bank $i \in \{1,2\}$ bids, and index by $s_i \in \{1,\ldots,S\}$ each ONM increment that occurs after bank i falls OUT at τ_{J_i} after making its final bid at ρ_{J_i}. Thus t_{s_i} is the time the s_i^{th} increment occurs, r_{s_i} is the ONM rate at t_{s_i}, and in the summations below when $s_i = 1$ we set $r_{s_i-1} = r_{\rho_{J_i}}$, the ONM rate the last time the i^{th} bank bids in the auction. For notational convenience we abbreviate $W_{b_{J_i}}$ with W_i for each bank $i \in \{1,2\}$.

Four mutually exclusive scenarios explain why neither bank bids again. In one scenario neither bank has the opportunity to rebid because its final bid wins. The final bid of each bank $i \in \{1,2\}$ at ρ_{J_i} implies $v_i > r_{\rho_{J_i}}$, and so the contribution of this component to the likelihood is:

$$\mathcal{L}_{W_1,W_2}\left(F_{V_1,V_2}\right) \equiv \Pr\left(v_1 > r_{\rho_{J_1}} \text{ and } v_2 > r_{\rho_{J_2}}\right) \equiv 1 - F_{V_1,V_2}(r_{\rho_{J_1}}, r_{\rho_{J_2}})$$

In another scenario the first bank wins, and the second bank loses. Then $v_1 > r_{\rho_{J_1}}$, but more can be inferred about the second bank. For example if $r_{s_2-1} \leq v_2 < r_{s_2}$ the second bank would bid again if given the opportunity to do so between τ_{J_2} and t_{s_2}. Summing the joint probabilities (v_1, v_2) over values of $v_2 > r_{\rho_{J_2}}$, weighted by probabilities to reflect not having the opportunity to bid gives:

$$\begin{aligned}\mathcal{L}_{W_1,\widetilde{W}_2}\left(F_{V_1,V_2}\right) \equiv & \sum_{s_2=1}^{S} \Pr\left(v_1 > r_{\rho_{J_1}} \text{ and } r_{s_2-1} \leq v_2 < r_{s_2}\right)\left[1 - G_{\tau_{J_2}}(t_{s_2})\right] \\ & + \Pr\left(v_1 > r_{\rho_{J_1}} \text{ and } v_2 > \overline{r}\right)\left[1 - G_{\tau_{J_2}}(T)\right]\end{aligned}$$

The expression for the probability that only the second bank of the two wins the auction is found by simply reversing the bank indices in the expression above. The likelihood component for the remaining scenario, both banks losing the auction, is given by the expression:

$$\begin{aligned}
\mathcal{L}_{\widetilde{W}_1,\widetilde{W}_2}(F_{V_1,V_2}) \equiv\ & \Pr(v_1 > \overline{r} \text{ and } v_2 > \overline{r})[1 - G_{\tau_{J_1}}(T)][1 - G_{\tau_{J_2}}(T)] \\
& + \sum_{s_1=1}^{S} \Pr(r_{s_1-1} \leq v_1 < r_{s_1} \text{ and } v_2 > \overline{r}) \left[1 - G_{\tau_{J_1}}(t_{s_1})\right] \left[1 - G_{\tau_{J_2}}(T)\right] \\
& + \sum_{s_2=1}^{S} \Pr(r_{s_2-1} \leq v_2 < r_{s_2} \text{ and } v_1 > \overline{r}) \left[1 - G_{\tau_{J_2}}(t_{s_2})\right] \left[1 - G_{\tau_{J_1}}(T)\right] \\
& + \sum_{s_1=1}^{S} \sum_{s_2=1}^{S} \Pr\begin{pmatrix} r_{s_1-1} \leq v_1 < r_{s_1} \\ \text{and } r_{s_2-1} \leq v_2 < r_{s_2} \end{pmatrix} \left[1 - G_{\tau_{J_1}}(t_{s_1})\right] \left[1 - G_{\tau_{J_2}}(t_{s_2})\right]
\end{aligned} \tag{C.3}$$

The four lines in $(C.3)$ are formed by partitioning $\left[r_{\rho_{J_i}}, \overline{v}\right]$ four ways, by $\left[r_{\rho_{J_i}}, \overline{r}\right)$ and $[\overline{r}, \overline{v}]$ for $i \in \{1, 2\}$. If both v_1 and v_2 exceed $\overline{r}$, the ONM at T, then the only way both banks lose is if neither has an opportunity to bid after τ_{J_1} and τ_{J_2}; the joint probability of these events is given by the right of the first line. Similarly in the second line if $v_2 > \overline{r}$ the second bank would bid in the time interval $[\tau_{J_2}, T]$ if it had an opportunity to do so, thus explaining the term $1 - G_{\tau_{J_2}}(T)$. If r_{s_1-1} and r_{s_1} bound v_1, then the first bank would bid in the interval $[\tau_{J_1}, t_{s_1})$, but not in the time interval $[t_{s_1}, T]$; thus $1 - G_{\tau_{J_1}}(t_{s_1})$ is the probability that the bank would not bid after τ_{J_2}. This explains the second line and, reversing the bank indices, the third. The rationale for both banks in the fourth line follows the same argument given for the first bank in the second line: the banks would decline opportunities to bid in the time intervals $[t_{s_1}, T]$ and $[t_{s_2}, T]$.

The contribution of a bank pair to the likelihood of observing any given win/loss combination is thus:

$$\begin{aligned}\mathcal{L}\left(F_{V_1,V_2}\right) &= \mathbb{1}\{W_1,W_2\}\mathcal{L}_{W_1,W_2}\left(F_{V_1,V_2}\right)+\mathbb{1}\{W_1,\widetilde{W}_2\}\mathcal{L}_{W_1,\widetilde{W}_2}\left(F_{V_1,V_2}\right) \\ &\quad +\mathbb{1}\{\widetilde{W}_1,W_2\}\mathcal{L}_{\widetilde{W}_1,W_2}\left(F_{V_1,V_2}\right)+\mathbb{1}\{\widetilde{W}_1,\widetilde{W}_2\}\mathcal{L}_{\widetilde{W}_1,\widetilde{W}_2}\left(F_{V_1,V_2}\right)\end{aligned} \tag{C.4}$$

C.3 Concavity of the likelihood

To establish the concavity of the likelihood function, let $\nu_1,\nu_2,...,\nu_M$ denote the grid of values which is a subset of the observed ONM rates, and let p_m for $m=1,...,M$ denote the probability that $v_i=\nu_m$. Following the notation in the main text we denote by $s=0,...,S$ the increases to the ONM rate, so that t_s is the time of the s^{th} ONM increase, r_s is the s^{th} ONM rate, and ρ denotes the time of the last bid placed by a bank (where $\rho>T$ if the bank finishes OUT). The likelihood for a single bank's bids is:

$$\begin{aligned} l &= \sum_{s=0}^{S}\Bigg\{\mathbb{1}\{t_s<\rho<t_{s+1}\}\left[1-\sum_{i=1}^{M}\sum_{m=1}^{i}\mathbb{1}\{r^{(i-1)}<r_s<r^{(i)}\}p_m\right] \\ &\quad +\mathbb{1}\{\rho>T\}(G(t_{s+1})-G(t_s))\sum_{i=1}^{M}\sum_{m=1}^{i}\mathbb{1}\{r^{(i-1)}<r_s<r^{(i)}\}p_m\Bigg\}+(1-G(T)) \\ &= \sum_{s=0}^{S}\mathbb{1}\{t_s<\rho\le t_{s+1}\}+(1-G(T)) \\ &\quad +\sum_{s=0}^{S}\sum_{i=1}^{M}\sum_{m=1}^{i}\mathbb{1}\{r^{(i-1)}<r_s<r^{(i)}\}\left(\mathbb{1}\{\rho>T\}(G(t_{s+1}-G(t_s))-\mathbb{1}\{t_s<\rho\le t_{s+1}\}\right)p_m \\ &= \sum_{s=0}^{S}\mathbb{1}\{t_s<\rho\le t_{s+1}\}+(1-G(T)) \\ &\quad +\sum_{m=1}^{M}\left[\sum_{i=m}^{M}\left(\ \left\{\textstyle\sum_{s=0}^{S}\left[\mathbb{1}\{\rho>T\}(G(t_{s+1}-G(t_s))-\mathbb{1}\{t_s<\rho\le t_{s+1}\}\right]\right\}\ \right)\right]p_m \end{aligned}$$

Define:

$$A \equiv (1 - G(T)) + \sum_{s=0}^{S} \mathbb{1}\{t_s < \rho \leq t_{s+1}\},$$

$$h^{(m)} \equiv \sum_{i=m}^{M} \left(\ \left\{ \sum_{s=0}^{S} \left[\mathbb{1}\{\rho > T\}(G(t_{s+1} - G(t_s)) - \mathbb{1}\{t_s < \rho \leq t_{s+1}\} \right] \right\} \ \right).$$

Then the likelihood can be expressed:

$$l = A + \sum_{m=1}^{M} h^{(m)} p_m.$$

We need only to consider the second term when maximizing the likelihood. The structural log-likelihood is thus:

$$\sum_{m=1}^{M} \ln(h^{(m)} p_m).$$

Noting that $p_M = 1 - \sum_{m=1}^{M-1} p_m$, the first order condition with respect to k is:

$$\frac{h^{(k)} - h^{(M)}}{\sum_{m=1}^{M} h^{(m)} p_m} = 0. \tag{C.5}$$

To establish concavity, we show that the Hessian H is negative definite by proving that the quadratic form $Q(x) = x^T H x$ is negative. Taking the cross-partial of (C.5) with respect to k' yields:

$$H_{kk'} = \frac{-(h^{(k)} - h^{(M)})(h^{(k')} - h^{(M)})}{(\sum_{m=1}^{M} h^{(m)} p_m)^2}.$$

Thus the quadratic form is:

$$\begin{aligned} Q(x) &= x^T H x \\ &= \sum_{k=1}^{M}\sum_{k'=1}^{M} x_k H_{kk'} x_{k'} \\ &= -\frac{\sum_{k=1}^{M}\sum_{k'=1}^{M} x_k (h^{(k)} - h^{(M)})(h^{(k')} - h^{(M)}) x_{k'}}{(\sum_{m=1}^{M} h^{(m)} p_m)^2} \\ &= -\left(\frac{\sum_{k=1}^{M} x_k (h^{(k)} - h^{(M)})}{\sum_{m=1}^{M} h^{(m)} p_m}\right)^2 < 0 \end{aligned}$$

Therefore the log-likelihood is concave because the Hessian is negative definite. Applying the same approach shows that the joint log-likelihood for valuation pairs has the same structure, linear in the joint probabilities, and so is concave too.

D Estimation

This section provides a consistent estimator of the identified portions of the valuation distribution, before giving details about its implementation and explains the subsampling procedure used to obtain the estimated standard errors.

D.1 A consistent estimator

Theorem 1 *Let $\mathcal{V}$ denote the set of ONM rates that have strictly positive density or mass on $[0, \overline{v}]$, and denote by $F^*(v)$ the cumulative distribution function that $\mathcal{V}$ induces on $F(v)$.*[18] *Denote by P_N a partition of $[0, \overline{v}]$ formed from $M_N - 1$ connected intervals indexed by $\left\{x_m^{(N)}\right\}_{m=1}^{M_N}$, where $0 \equiv x_1^{(N)} < \ldots < x_{M_N}^{(N)} \equiv \overline{v}$ for each $N = 1, 2, \ldots$. We impose the following conditions: (i) $M_N^{-1} = o(1)$; (ii) $M_N = o(N)$; (iii) $\max\left\{x_{m+1}^{(N)} - x_m^{(N)} : m = 1, \ldots M_N - 1\right\} = o(1)$*

[18] Thus $F^*(v) = F(v)$ for all $v \in \mathcal{V}$ and is nonincreasing for all $v \notin \mathcal{V}$.

and (iv) $P_N \subseteq P_{N+1} \subseteq \ldots$. *Denote by* $F_N^*(v)$ *the step function distribution induced by* $F^*(v)$ *on to* P_N*, and denote the maximum likelihood estimator for* $F_N^*(v)$ *by* $\hat{F}_N^*(v)$*. Then* $\sup_{v\in\mathcal{V}} |\hat{F}_N^*(v) - F^*(v)| \overset{a.s.}{\to} 0$*.*

Proof. The proof adapts the Glivenko-Cantelli theorem to our model. Noting that both $F^*(v)$ and $\hat{F}_N^*(v)$ are monotone increasing in $v \in [\underline{v}, \overline{v}]$ it follows that for all $x_m^{(N)} \leq v \leq x_{m+1}^{(N)}$:

$$\hat{F}_N^*(x_m^{(N)}) \leq \hat{F}_N^*(v) \leq \hat{F}_N(x_{m+1}^{(N)}) \text{ and } F^*(x_m^{(N)}) \leq F^*(v) \leq F^*(x_{m+1}^{(N)})$$

Together these inequalities imply:

$$\begin{aligned}
\hat{F}_N^*(v) - F^*(v) &\leq \hat{F}_N^*(x_{m+1}^{(N)}) - F^*(x_m^{(N)}) - F^*(x_{m+1}^{(N)}) + F^*(x_m^{(N)}) \\
&= \left[\hat{F}_N^*(x_{m+1}^{(N)}) - F^*(x_{m+1}^{(N)})\right] + \left[F^*(x_{m+1}^{(N)}) - F^*(x_m^{(N)})\right] \\
\hat{F}_N^*(v) - F^*(v) &\geq \hat{F}_N^*(x_m^{(N)}) - F^*(x_{m+1}^{(N)}) + F^*(x_m^{(N)}) - F^*(x_m^{(N)}) \\
&= \left[\hat{F}_N^*(x_m^{(N)}) - F^*(x_m^{(N)})\right] + \left[F^*(x_m^{(N)}) - F^*(x_{m+1}^{(N)})\right] \qquad \text{(D.1)}
\end{aligned}$$

Fix any $\epsilon > 0$, and choose N so that $F^*(x_{m+1}^{(N)}) - F^*(x_m^{(N)}) < \epsilon$ for all $x_m^{(N)} \in P_N$. Note this is possible by conditions (iii) and since $F^*(v)$ is right continuous. Then (D.1) implies:

$$\left[\hat{F}_N^*(x_m^{(N)}) - F^*(x_m^{(N)})\right] - \epsilon \leq \hat{F}_N^*(v) - F^*(v) \leq \left[\hat{F}_N^*(x_{m+1}^{(N)}) - F^*(x_{m+1}^{(N)})\right] + \epsilon$$

Appealing to (iv), the strong law of large numbers, and the concavity of the likelihood function, $\hat{F}_N^*(x) \overset{a.s.}{\to} F^*(x)$ for all $x \in \left\{x_m^{(N)}\right\}_{m=1}^{M_N}$. Hence $\hat{F}_N^*(v) \overset{a.s.}{\to} F^*(v)$ for all $v \in \mathcal{V}$, proving the theorem. ■

D.2 Implementation

The criterion function (D.2) used to estimate $F_V(\cdot)$ is constructed from the individual bank's likelihood contribution (4) by taking the product over all auctions $k \in \{1, ..., K\}$ and banks $i \in \{1, ..., I\}$:

$$\mathcal{L}(F_V, \{\hat{G}_t\}; \{\tau_{J_{ik}}, \rho_{J_{ik}}\}) = \prod_{k=1}^{K} \prod_{i=1}^{I} \left([1 - F_V(r_{\rho_{J_{ik}}})] + \mathbb{1}\{\widetilde{W}_{b_{J_{ik}}}\} H(F_V, \hat{G}_{\tau_{J_{ik}}}) \right) \tag{D.2}$$

Note (D.2) would be a likelihood function for $F_V(\cdot)$ if the true $G_t(\cdot)$ were used instead of a first-stage estimate $\hat{G}_t(\cdot)$.

To estimate the joint distribution $F_{V_1,V_2}(\cdot)$ we construct a criterion function similar to (D.2). The key difference is that pairs of banks are matched up in taking the product over auctions. Because the number of distinct pairs in an auction with N competing banks is $N!$ we randomly select 200 bank pairs per auction for each, form a discrete approximation to F_{V_1,V_2} which we denote by $F^*_{V_1,V_2}$, and maximize a criterion function over the set of quantiles for the discretized joint distribution:

$$\hat{F}_{V_1,V_2} = \arg \max_{F^*_{V_1,V_2}} \prod_{k \in K} \prod_{(i,j) \in I_k} \mathcal{L}\left(F^*_{V_1,V_2}, \{\hat{G}_t\}; \{\tau_{J_{ik}}, \rho_{J_{ik}}\} \right) \tag{D.3}$$

where K is the set of auctions and I_k is the set of randomly sampled pairs of banks (i, j) in auction k.

D.3 Subsampling algorithm

To estimate confidence intervals we follow the subsampling approach of Politis, Romano, and Wolf (1999). A standard subsampling algorithm uses subsamples of size $b << A$ to re-estimate the parameters of interest, where A is the number of auctions. Denote our parameter estimate by $\widehat{\theta}_A$ and the parameter estimate from subsample a of size b by $\widehat{\theta}_{A,b,a}$. We create q subsamples and re-estimate the parameter on these subsets using the following as an approximation for the sampling distribution:

$$L_{A,b,||.||}(x|\tau_b) \equiv q^{-1} \sum_{a=1}^{q} \mathbf{1} \left\{ \tau_b ||\widehat{\theta}_{n,b,a} - \widehat{\theta}_n|| \leq x \right\}$$

where τ_b is the convergence rate of the estimator, q represents the number of all possible $\begin{pmatrix} A \\ b \end{pmatrix}$ combinations of the data, and $|| \cdot ||$ is the Euclidean norm for the linear space Θ. Theorem 2.5.2 on page 53 of Politis, Romano, and Wolf (1999) establishes the asymptotic validity of this approach. They also prove that randomly selecting observations with or without replacement from the set of auctions $\{1, \ldots, A\}$ yields consistent estimates. This stochastic approximation substantially reduces the computational burden relative to choosing all possible combinations. The confidence intervals in Section 6 are generated from 400 subsamples. Because of within-auction correlation induced by the unobserved heterogeneity term, we sample entire auctions rather than individual bids. Thus each subsample is created from all the bids in 20 auctions randomly chosen from the full sample.