Auctions in Financial Markets*

Jakub Kastl†

August 19, 2019

In this article I discuss how auctions and tools developed for their empirical analysis can inform empirical analysis of financial markets. Since virtually all markets organized as auctions have well-specified and known rules that map nicely into game-theoretical models, I demonstrate using several applications that one can often leverage particular details to study issues that have nothing to do with the auction per se. To do so, I first review an estimation method, which is widely applicable in many settings where a researcher needs to recover agents' beliefs, in order to establish a link between observables and unobservables using some version of a necessary condition for optimality. I then discuss applications to quantification of front-running, evaluation of quantitative easing operations and estimation of a demand system for financial products.

*This paper is based on a keynote lecture given at 2018 EARIE Meetings in Athens. Thanks to many co-authors for continuing collaboration and to Paul Heidhues, Xavier Vives and many colleagues for inspiration and feedback. The usual disclaimer applies.

†Department of Economics, Princeton University, NBER and CEPR

1 Introduction

The main goal of this article is to show how the tools recently developed in the Industrial Organization literature to analyze auction markets can be utilized to address questions that often arise in financial markets. The main theme throughout is that often particular details of how the various auctions are run provide us with an opportunity to look at the markets in a different and novel way. This paper complements an earlier article (Kastl (2017)) that also discusses the applications using IO tools in financial markets, but the focus here will be particularly on leveraging details of the auction mechanisms.

For example, a seemingly innocuous detail - that auctions of various maturities of government securities are frequently sold simultaneously and virtually the same bidders participate in all of these - allows us to sidestep one important problem typically encountered during estimation. Since we are able to observe how the willingnesses to pay for various maturities move together within a bidder and we are essentially able to observe a vector of those at the same time, we can study the linkages between the willingnesses to pay for different securities and recover the underlying demand system while not being exposed to issues involving unobserved heterogeneity, i.e., that other features of the environment might be changing and affecting the marginal valuation in an unobservable way. Another interesting detail from many Treasury bill auctions is that bids are often recorded with exact timestamps and they may also be revised before the auction deadline. Since some large participants, such as pension funds, are either required or choose to route their bids through primary dealers, we can utilize this observed timing of the bids and their dynamics to quantify (a particular version of) front-running. Finally, in the third application I will discuss how one can use the auction data to potentially analyze externalities across implementations of unconventional monetary policies when central banks do not coordinate their actions.

Theoretical analysis of auctions can be traced back at least to Vickrey (1961). The seminal article of Milgrom and Weber (1982) that methodically laid out equilibrium analysis of different auction formats in different information environments really opened the door for empirical analysis, both model-based and descriptive. Hendricks and Porter (1988) is an early empirical examination of oil lease auctions which provides a convincing argument that the behavior of actual bidders is close to what the Milgrom and Weber (1982) theory predicts. Laffont and Vuong (1996) and Guerre et al. (2000) then provide the key building block for structural analysis of auction data: the observation that bids can be inverted into values using the necessary condition for optimal bidding. This approach is reminiscent of the (Bresnahan, 1989) approach to identifying (unobserved) marginal cost from equations governing optimal pricing behavior within a differentiated Bertrand oligopoly.

In finance, the market microstructure foundational papers are Glosten and Milgrom (1985) and Kyle (1985, 1989). There are many papers that try to estimate the price impact of individual market participants based on Kyle's model. Shleifer (1986) uses stocks that are being added (or withdrawn) from a market-wide index - and the thus created shifter in demand due to the behavior of index funds - to show that the demand curve for US equities is downward-sloping. This implies that there is potential for price-impact of large orders. There has been a burgeoning (theoretical) recent literature on market microstructure, including Vives (2011) or Rostek and Weretka (2012). There has been little work taking these models to the data though. One important exception includes primary markets for government debt, because these are often organized as auctions.

Unlike many, if not most, other empirical settings, auctions have a major advantage. The rules that the market participants need to follow are known pretty much to the last detail and are often very close to how we write down theoretical models. Since we have the necessary theories that describe optimal behavior in various settings that frequently occur, we can thus use the appropriate models to learn about the models' fundamentals, such as the distribution of underlying values that are unobserved to the econometrician. This typically proceeds in two steps: for any action that we observe, we aim to use the rule for optimal behavior from the model to establish a link between an observable variable (e.g., the bid) and the unobservable that we need to recover (e.g., the marginal value). This rule, however, typically includes various expectations: the bidders need to optimize against actions of their rivals that depend on realizations on various variables that are unobserved to their rivals. Hence, the researcher needs to estimate the beliefs the bidders posses in order to make progress. Under rational expectations one can use the distribution of realized actions (typically together with some assumptions on the dependence of the random unobservables) to approximate the distribution of rivals' actions against which a bidder needs to optimize. Under stronger assumptions (such as iid), one can in fact augment the observed realized actions via simulations in order to approximate the beliefs without having to pool data across auctions, which helps alleviating potential issues with (auction-level) unobserved heterogeneity.

This paper proceeds as follows. In Section 2 I go through the typical model and how to approach estimation. Athey and Haile (2002) provides a similar general setup for single-unit auction environments. Since this paper is aimed at applications in financial markets I start with a theoretical model of a general share auction that yields a set of necessary conditions for optimal bidding which link the observables (bids) to the unobservables (marginal values). Since the applications in this paper will involve only the discriminatory (or pay-your-bid) auction format, I will focus solely on this format. Since for any particular auction rules these optimality conditions involve beliefs about rivals' actions, I review the resampling approach originally proposed in Hortaçsu (2002) that allows for easy estimation of these beliefs using observed behavior. Then I go through few applications from financial markets which utilize the details of the auction mechanisms to address various questions. First, in Section 3 I discuss two applications studying Canadian Treasury auctions. In subsection 3.2 I show how to quantify a particular version of front-running using exact timestamps from bids submitted by primary dealers in auctions from Canadian Treasury bills. Second, in subsection 3.3 I illustrate, using the data from Canadian Treasury bills as well, how to estimate a whole demand system for (a subset of) government securities utilizing timing of the auctions themselves. Finally, in Section 4 I discuss how data from QE auctions run in the U.K. can help us study the impact of unconventional monetary policy on secondary markets. Section 5 concludes.

2 How to Use Auction Data

I begin by discussing the main model used throughout this paper. It is based on a model of the share auction proposed by Wilson (1979).

Optimal behavior in any auction mechanism, whether for a single item or for multiple items, can be typically characterized (perhaps only implicitly) through an equation expressing the equilibrium bid roughly as willingness-to-pay minus a shading factor:

$$BID = WTP - SHADING. \tag{1}$$

The "SHADING" term, might be positive (as in a first price auction in an independent private value (or IPV) setting), zero (as in a second price auction in a IPV setting) or negative (as in a third price auction in a IPV setting) - depending on the auction rules. The strategic considerations players face in the game, which are a function of the rules of the mechanism and of the environment, thus make it hard to go from the observed data (bids) to the parameters of interest (values).

Hence, to proceed we will make use of the rules of the auction mechanism and the equilibrium assumption to express the "SHADING" terms as a function of objects that are estimable given the available data.

More generally, a theoretical model of how agents behave typically allows the researcher to express the unobservable objects of interest as some function (or correspondence) of the observables. Jovanovic (1989) provides a formalization. The basic idea is to determine under what conditions is this function one-to-one: whether we can (uniquely) go back and forth from observables to primitives and vice versa. After we obtain estimates of the WTP, we can start performing counterfactuals of interest. For example, we can estimate how much surplus a given auction mechanism fails to extract or how inefficient an allocation it implements. Together, these two quantities bound how much revenue is lost. Ausubel et al. (2014) demonstrate the impossibility to rank the different auction mechanisms based on either efficiency or revenue theoretically. This implies that empirical analysis of each case is necessary.

2.1 Model of a Discriminatory Share Auction

The starting point is a share auction model based on Wilson (1979). Kastl (2012) extends this model to setting, in which bidders are restricted to employ step-function strategies. I start with the basic symmetric model with private information and private values and will introduce its variants and various generalizations as we move on to the applications. I discuss this general model in similar detail in Kastl (2017) and its variants appear in the papers cited therein. As an alternative to this model Kyle (1989), Vives (2011), Rostek and Weretka (2012) and Ausubel et al. (2014) use special functional form assumptions on payoffs and distributions of random variables instead. This allows them to obtain closed-form solutions for equilibrium demands, which allows them to analyze various comparative statics. Instead, I will focus here on implicit equilibrium characterization via a set of necessary conditions for optimal bidding. Let N be the number of potential bidders and S_i be an $D \geq 1$ dimensional private signal that i observes. This signal affects the underlying value for the auctioned good.

Assumption 1 *Bidders' signals, $S_1, ..., S_N$, are drawn from a common support $[0,1]^D$ according to an atomless joint d.f. $F(\cdot)$ with density f.*

For the characterization results below some conditions on the smoothness of the distribution of residual supplies is needed. This can come either from properties of $F(\cdot)$ or one can assume uncertain continuously distributed supply as is common in electricity auctions (and which will also be required for consistency of one of the proposed estimators at least for one version of the asymptotic argument). I will for simplicity assume that there is exogenous uncertainty about the supply, Q.

Assumption 2 *Supply Q is a random variable distributed on $\left[\underline{Q}, \overline{Q}\right]$ with strictly positive density conditional on S_i $\forall i$.*

Obtaining a share q of the supply Q is valued according to a marginal valuation function $v_i(q, S_i)$. I will assume that this function satisfies some regularity conditions.

Assumption 3 *$v_i(q, S_i)$ is non-negative, bounded, strictly increasing in each component of S_i $\forall q$, and weakly decreasing and continuous in q $\forall S_i$.*

Note that under Assumption 3, values are assumed to be private since $v_i(\cdot)$ does not depend on private information of the rivals. Since in many settings (and particularly in financial markets) we may worry that values are interdependent and marginal value of bidder i might depend on signal of bidder j, it is important to provide more discussion of this assumption. Hortaçsu and Kastl (2012) uses the bid updating by primary dealers to test whether there are statistically significant information spillovers between customers and primary dealers. I will discuss this a bit more below.

Furthermore, as mentioned above to respect the rules governing the actual auction markets, I will restrict the strategy set available to each bidder to step functions with at most K steps. I will use the index k to denote the place of a particular bidpoint in the vector of bids that are ordered to be increasing in the quantity dimension and decreasing in price. A bidpoint (q_k, b_k) together with the preceding bidpoint (q_{k-1}, b_{k-1}) thus specify the marginal quantity, $q_k - q_{k-1}$, that a bidder is bidding b_k for. I will also allow a bid l, which is sure to lose, and hence basically corresponds to not participating in the auction.

Assumption 4 *Each player $i = 1, ..., N$ has an action set:*

$$A_i = \left\{ \begin{array}{c} \left(\vec{b}, \vec{q}, K_i\right) : \dim\left(\vec{b}\right) = \dim\left(\vec{q}\right) = K_i \in \{1, ..., K\}, \\ b_{ik} \in B = \{l\} \cup \left[0, \bar{b}\right], \; q_{ik} \in [0,1], \; b_{ik} > b_{ik+1}, q_{ik} < q_{ik+1} \end{array} \right\}.$$

Finally, I will assume that rationing is pro-rata on the margin. This means that if demand at the market clearing price were to exceed the supply, all bidders' bids submitted exactly at that price would be adjusted proportionally.

Henceforth, I will use the term K-step equilibrium to refer to a Bayesian Nash Equilibrium of the auction game where the bidders are restricted to use at most K steps. In other words, a K-step equilibrium is a profile of strategies, such that each strategy is a step function with at most K steps and it maximizes the expected payoff for (almost) every type s_i for all i.

2.1.1 Equilibrium Characterization

I begin by laying out the equilibrium of a discriminatory auction: an auction in which every bidder has to pay her bid for all units he/she wins. A version of this mechanism is used in all three applications discussed here. Kastl (2017) provides a similar discussion also for the case of a uniform price auction. Let $V_i(q, S_i)$ denote the gross utility: $V_i(q, S_i) = \int_0^q v_i(u, S_i)\, du$. The expected utility of a bidder i of type s_i employing a strategy $y_i(\cdot|s_i)$ can be written as:

$$\begin{aligned} EU(s_i) \quad = \quad & \sum_{k=1}^{K_i} \left[\Pr(b_{ik} > P^c > b_{ik+1}|s_i) V(q_{ik}, s_i) - \Pr(b_{ik} > P^c|s_i) b_{ik}(q_{ik} - q_{ik-1})\right] \qquad (2) \\ & + \sum_{k=1}^{K_i} \Pr(b_{ik} = P^c|s_i) E_{Q,s_{-i}|s_i}\left[V\left(Q_i^c(Q, \mathbf{S}, \mathbf{y}(\cdot|S)), s_i\right) - b_{ik}\left(Q_i^c(Q, \mathbf{S}, \mathbf{y}(\cdot|S)) - q_{ik-1}\right)|b_{ik} = P^c\right] \end{aligned}$$

where I let $q_{i0} = b_{iK_i+1} = 0$. The random variable $Q_i^c(Q, \mathbf{S}, \mathbf{y}(\cdot|S))$ is the (market clearing) quantity bidder i obtains if the state (bidders' private information and the supply quantity) is $(Q, \mathbf{S})$ and bidders submit bids specified in the vector $\mathbf{y}(\cdot|S) = [y_1(\cdot|S_1), ..., y_N(\cdot|S_N)]$. Since bidders potentially have private information and since the supply is random, the market clearing price is a random variable, denoted by P^c.

A Bayesian Nash Equilibrium in this setting is a collection of functions such that (almost) every type s_i of bidder i is choosing his bid function so as to maximize her expected utility: $y_i(\cdot|s_i) \in \arg\max EU_i(s_i)$ for a.e. s_i and all bidders i. The system of necessary conditions implicitly characterizing such a BNE is the link between the observables and unobservables that we seek to establish. Kastl (2012) derives the following necessary conditions that the quantity requested in any step of a pure strategy that is a part of a K-step equilibrium has to satisfy.

Proposition 1 *Under assumptions 1–4 in any K-step Equilibrium of a discriminatory auction,*

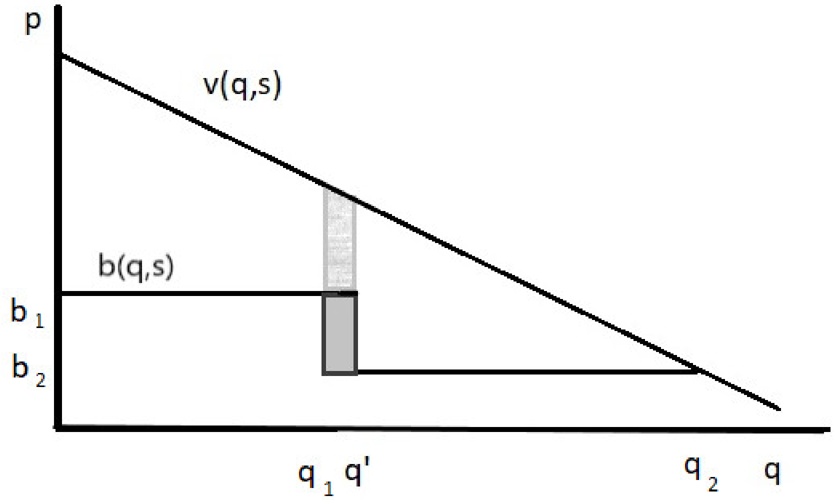

Figure 1: Trade-off in a Discriminatory Auction

for almost all s_i, every step $k < K_i$ in the equilibrium bid function $y_i(\cdot|s_i)$ has to satisfy

$$\Pr\left(b_{ik} > P^c > b_{ik+1}|s_i\right)\left[v\left(q_{ik}, s_i\right) - b_{ik}\right] = \Pr\left(b_{ik+1} \geq P^c|s_i\right)\left(b_{ik} - b_{ik+1}\right) \tag{3}$$

and at the last step K_i it has to satisfy $v\left(\overline{q}, s_i\right) = b_{iK_i}$ where $\overline{q} = \sup_{q,s_{-i}} Q_i^c\left(q, s_{-i}, s_i, \mathbf{y}\left(\cdot|S\right)\right)$.

Note that this condition is simply a multi-unit counterpart of the equilibrium condition for bidding in a first-price auction: $g(b)\left(v - b\right) = G(b)$, where $G(b)$ is the CDF (and $g(b)$ is the PDF) of the distribution of the first-order statistic of rival bids (i.e., of the highest of rival bids) (Guerre et al., 2000). The trade-off in the multi-unit environment remains virtually the same. The bidder is simply trading off the expected surplus on the marginal (infinitesimal) unit versus the probability of winning it. Figure 1 illustrates this trade-off. When contemplating a deviation from a given bid characterized by two steps, $\{(b_1, q_1), (b_2, q_2)\}$, to one with q_1 replaced by q', the bidder of type s is essentially comparing the loss of surplus on the extra units, $\Delta q * (v(q_1, s) - b_1)$, captured by the light grey-shaded trapezoid, against the gain in surplus on that unit, if the market clearing price were to be even lower than b_2 - depicted by the dark grey-shaded rectangle, each weighted by their respective probability. The loss occurs whenever the market clearing price ends up being in the interval (b_2, b_1), i.e., with probability $\Pr(b_1 > P^c > b_2)$, in which case this bidder loses this marginal unit by moving it from the high bid, b_1, to the low bid, b_2 and it thus goes unfilled. The gain occurs whenever the market clearing price falls weakly short of b_2, i.e., with probability $\Pr(b_2 \geq P^c)$.

Using (3) we can go from observables (bids) to the willingness-to-pay $v\left(q, s_i\right)$. The estimation approach I will utilize is similar in spirit to Laffont and Vuong (1996) and Guerre et al. (2000) and has been proposed for the multi-unit auction environment first by Hortaçsu (2002). Note that (3) is simply a necessary condition for an optimal choice of q_k. The set of these conditions (one for each k) thus identifies K points of the function $v(q, s_i)$. The optimality condition with respect to the bid b_k can be derived in a straightforward manner by differentiating (2), but it cannot be simplified and interpreted as naturally as equation (3). Nevertheless, it may allow a researcher to obtain tighter bounds on the function $v(q, s_i)$. See McAdams (2008) for more general partial identification results.

The question of equilibrium existence in the environment where bidders bid using step functions has been addressed in Kastl (2012). Proposition 2 in that paper establishes existence of an equilibrium in distributional strategies for a discriminatory auction with private values and supply uncertainty whenever the potential dependence between private signals is not "too large." Formally, it is required that the probability measure associated with the distribution function $F(S_1, ..., S_N)$ is absolutely continuous with respect to $\prod_{i=1}^{N} \mu_i$ where μ_i is the marginal distribution of S_i. For the case of symmetrically informed bidders, but uncertain supply, Pycia and Woodward (2015) show that there is a unique pure-strategy Baysian-Nash equilibrium.

2.2 Estimation

The estimation approach described in what follows can be generalized to many other settings than just auctions. The main point is that the relationship between observables (bids in case of an auction, submitted student preferences in case of school matching mechanisms, buy–sell orders in case of data from equity markets etc) and objects of interest, i.e., preferences, valuations or willingness-to-pay, is given by an equilibrium relationship of an explicit economic model, in which each participant's behavior depends on her beliefs about rivals' behavior. Typically the researcher assumes a Nash Equilibrium (or its appropriate form such as a Bayesian Nash Equilibrium or a refinement such as a Perfect Bayesian Equilibrium when selection is necessary) which then restricts participants' beliefs about rivals' play to be consistent with the actual strategies. In many environments or games, the equilibrium strategies depend on the rivals' actions only indirectly, and the actual uncertainty is about some functional of these actions. For example, in the applications discussed here, the uncertainty is about where the market will clear, i.e., about the market clearing price. In case of school choice, the uncertainty is about whether a given school might have an open spot for the applicant. In case of equity markets, the uncertainty might be about whether a counterparty for a particular trade can be found. One way to account for this uncertainty in estimation non-parametrically is to assume that all participants agree on its distribution ex-ante and hence any differences in their observed strategies ex-post are due to differences in realizations of their private information. In that case, one can account for the uncertainty about the object of interest and thus estimate its probability distribution by "resampling," i.e., by following a bootstrap-like procedure where by drawing repeated samples of strategies (by sampling with replacement from the observed data) one simulates different possible states of the world and thus eventually obtains an estimate of the distribution of the random variable(s) of interest. This idea appeared originally in Hortaçsu (2002) and was later applied in Hortaçsu and McAdams (2010) and Kastl (2011). Such an estimator is a particular form of a V-statistic, which has useful implications for characterizing its asymptotic behavior. It should be intuitive to see that as one constructs more and more samples (as the data set gets larger and thus the observed strategies span more and more of the type space), one obtains more and more simulated states of the world, where the probability of a type profile being in some subset of the type space $[0,1]^M \times [\underline{Q}, \overline{Q}]$ corresponds to the probability of that subset implied by the population distribution functions given in Assumptions 1 and 2. I will now describe this approach in the context of auctions in detail.

2.2.1 Typical Data Set

A typical auction data set $\{B_{1t}, ..., B_{Nt}, Q_t, X_t\}_{t=1}^T$, consists of T auctions and for each t, the set of bids $B_{1t}, ..., B_{Nt}$, the realized supply, Q_t, and some auction covariates, X_t, where N is the number of (potential) bidders. In the auctions we will talk about here, B_{it} is typically a function $b_{it}(q)$ which maps quantity requested (e.g., the face value of Treasury bills the bidder is bidding for) into the bid, typically specified in terms of yield. In most applications, $b_{it}(\cdot)$ takes form of the step function, $\{(q_{itk}, b_{itk})\}_{k=1}^{K_{it}}$ where K_{it} number of steps bidder i uses in auction t. Supply may be deterministic and announced by the auctioneer prior to the auction (as is the case with US Treasury auctions) or may be stochastic since the auctioneer may adjust the preannounced quantity later (as is the case with the ECB auctions of short-term loans). In some applications, bidders might be of different types, which may either be observable or latent. For example, we might have primary dealers and other bidders and we may have a good reason to believe that the underlying distribution of their willingness-to-pay may differ across these groups. We may therefore have C bidder classes, and thus have N_{ct} potential bidders of class c in auction t.

2.2.2 Estimating the Distribution of the Market Clearing Price

Equation (3) reveals that to get at the object of interest, $v(\cdot)$, we need to estimate the distribution of P^c, $H(X) = \Pr(P^c \leq X)$. With such an estimate in hand, one can simply plug-in to (3) and obtain an estimate of $v(\cdot)$. $H(X)$ is defined as

$$H(X) \equiv \Pr(X \geq P^c | s_i) = E_{\{Q, S_{j \neq i}\}} I \left(Q - \sum_{j \neq i} y(X|S_j) \geq y(X|s_i) \right) \tag{4}$$

where $I(\cdot)$ is the indicator function and $E_{\{.\}}$ is an expectation over the random supply and other bidders' private information. The market clearing price will be lower than X whenever the realized types and the realization of the supply are such that the implied realized aggregate residual supply (i.e., total supply minus demands of everyone but i) falls short of i's demand. Hortaçsu and Kastl (2012) describes how to estimate $H(\cdot)$. The approach begins by defining an indicator of excess supply at price X (given bid functions $\{y_j(X|s_j)\}_{j\neq i}$ and i's own bid $y_i(X|s_i)$) as follows:

$$\Phi\left(\{y_j(X|s_j)\}_{j\neq i}; X\right) = I\left(Q - \sum_{j\neq i} y_j(X|s_j) \geq y_i(X|s_i)\right).$$

One estimator of $H(X)$ can be derived as a V-statistic:

$$\xi\left(\hat{F}; X, h_T\right) = \frac{1}{(NT)^{(N-1)}} \sum_{\alpha_1=(1,1)}^{(T,N)} \ldots \sum_{\alpha_{N-1}=(1,1)}^{(T,N)} \Phi\left(y_{\alpha_1}, \ldots, y_{\alpha_{N-1}}, X\right) \tag{5}$$

where $\alpha_i \in \{(1,1),(1,2),\ldots(1,N),\ldots,(T,N)\}$ is the index of the bid in the subsample and $\hat{F}$ is the empirical distribution of bids, i.e., the empirical probability distribution over points in 2K-dimensional space.

It is straightforward to see that an estimator defined as a simulator of $\xi\left(\hat{F}; X, h_T\right)$ by drawing only M subsamples rather than all $(NT)^{(N-1)}$ is consistent as $T \to \infty$ (and under appropriate conditions on the rate at which the number of simulations, M, increases – see Pakes and Pollard (1989) for more details). Cassola et al. (2013) establish that it is also consistent as $N \to \infty$ provided some further technical conditions are satisfied (one of which is that there is a non-vanishing uncertainty about supply per bidder in the limit). This latter asymptotic result is particularly important when studying financial markets and auctions in general (but also in other applications), since it allows to use data from one auction at a time and thereby keep constant many unobservable characteristics. When the economic environment is changing and it is hard to capture these changes either through covariates or the data is not rich enough for the researcher to be able to employ some sophisticated econometric technique that would control for the unobserved heterogeneity in the spirit of Krasnokutskaya (2011) a researcher might not be comfortable pooling data across auctions. In such an environment, thinking about asymptotics in N within a fixed auction might be preferred. Hortaçsu and Kastl (2012) further show how to modify this estimator in order to allow for asymmetries and the presence of covariates by introducing weighting into the resampling procedure which is similar to a nonparametric regression.

Applying the estimation approach described above, we can thus obtain an estimate of the "Shading" factor in (1) that determines the wedge between the submitted bid and the actual willingness-to-pay. With an estimate of bidder's actual willingness-to-pay, which is of course ultimately equivalent to having an estimate of that bidder's demand curve and we can start evaluating how the expected utility associated with these demand curves changes as information from a customer arrives.

3 Canadian Treasury Bill Auctions

3.1 Primer on the Canadian Treasury Market

The Canadian Treasury finances new debt and re-finances old debt via a regular discriminatory auction. Typically, the auctioneer announces the size of the issuance, bidders submit (sealed) bids, currently via an electronic interface, and the market clearing price is determined as the highest price (or lowest yield) at which demand equals supply. There are two types of bids that bidders may submit: (i) noncompetitive (up to a limit of $5 million), whereby a bidder guarantees himself an allocation of the requested amount at a price that is equal to the quantity-weighted average accepted bid and (ii) competitive, in which both the price and quantity can be freely specified. The market for new issuance of debt has been historically organized around a fairly small number of "primary dealers." In addition to primary dealers, there are also other participants whose bids the Treasury wants separately identified such as large pension funds or banks that choose not to become primary dealers. I will denote such bidders simply as "customers" since they have to submit their bids through one of the primary dealers.

3.2 Quantification of Front-running

3.2.1 Modification of the Model of Bidding

We now need to modify the model of bidding introduced earlier to explicitly incorporate that there are two different groups of bidders, primary dealers and customers. As mentioned previously, the model, and the measurement it will allow us to conduct, will *rely on the assumption of bidder optimization.* In order to do that I modify Assumption 1 as follows:

Assumption 1′ *Customers' and dealers' private signals are independent and drawn from a common support* $[0,1]^D$ *according to atomless distribution functions* $F^P(.)$ *and* $F^C(.)$ *with strictly positive densities.*

Similarly, Assumption 3 is replaced by one allowing for a heterogeneous valuation function: $v^g\left(q, S_i^g\right)$ for $g \in \{P, C\}$ for "Primary," and "Customer." Note that, as mentioned before, Assumption 3 (and of course even its current modification) implies that learning other bidders' signals does not affect one's own valuation – i.e., we have a setting with *private*, not interdependent values. This assumption may be more palatable for certain securities (such as shorter-term securities, which are essentially cash substitutes) than others. In Section 3.2.3 I will describe a formal test of this assumption for the case of the Canadian Treasury bill market, and in that context private values cannot be rejected. Note that under this assumption, the additional information that a primary dealer j possesses due to observing her customers' orders, Z_j^P, simply consists of those submitted orders. As will become clear below, this extra piece of information allows the primary dealer to update her beliefs about the competitiveness of the auction, or, somewhat more precisely, the distribution of the market clearing price.

A primary dealer will condition on whatever information is observed in her customers' bids. In a (conditionally) independent private values environment, this information does not affect the primary dealer's own valuation, or her inference about other bidders' valuations, it only impacts her beliefs about the location and shape of the residual supply. Let $\mathcal{C}, \mathcal{P}$ denote the index sets containing indexes of customers and primary dealers, respectively. The distribution of the market clearing price from the perspective of primary dealer j, who observes the bids submitted by customers contained in an index set $\mathcal{C}_j$, is given by:

$$\begin{aligned}&\Pr\left(p \geq P^c|s_j, z_j\right) = \\ &\qquad E_{\left\{S_{k\in\mathcal{C}\backslash\mathcal{C}_j}, S_{n\in\mathcal{P}\backslash j}, Z_{n\in\mathcal{P}\backslash j}|z_j\right\}} I\left(RS\left(p, Q, \vec{S}, \vec{Z}\right) \geq y^P\left(p|s_j, z_j\right) + \sum_{m\in\mathcal{C}_j} y^C\left(p|s_m\right)\right)\end{aligned} \tag{6}$$

where $RS\left(p, Q, \vec{S}, \vec{Z}\right) = Q - \sum_{k\in\mathcal{C}\backslash\mathcal{C}_j} y^C\left(p|S_k\right) - \sum_{n\in\mathcal{P}\backslash j} y^P\left(p|S_n, Z_n\right)$, i.e., the residual supply at price p given supply realization Q and realization of private information $\left(\vec{S}, \vec{Z}\right)$. This expression simply says that the dealer "learns about competition" – the primary dealer's expectations about the distribution of the market clearing price are altered once she observes a customer's bid.

Finally, the distribution of P^c from the perspective of a customer is very similar to an uninformed primary dealer, but with the additional twist that the indirect bidder recognizes that her bid will be observed by a primary dealer, d, and can condition on the information that she provides to this dealer. The distribution of the market clearing price from the perspective of an indirect bidder j, who submits her bid through a primary dealer d is given by:

$$\begin{aligned}&\Pr\left(p \geq P^c|s_j\right) = \\ &\qquad E_{\left\{S_{k\in\mathcal{C}\backslash j}, S_{n\in\mathcal{P}}, Z_{n\in\mathcal{P}}|s_j\right\}} I\left(Q - \sum_{k\in\mathcal{C}\backslash j} y^C\left(p|S_k\right) - \sum_{n\in\mathcal{P}} y^P\left(p|S_n, Z_n\right) \geq y^C\left(p|s_j\right)\right)\end{aligned} \tag{7}$$

where $y^C\left(p|s_j\right) \in Z_d$ and $d \in \mathcal{P}$.

Note that the probability distributions from perspective of an uninformed primary dealer, or that of an informed primary dealer (given her observation of customer order flow) and that of an indirect bidder can then be estimated using equations (6), and (7) using the resampling technique described above. With the estimates of the probability distributions of market clearing price in hand, we can use equation (3) to estimate the willingness-to-pay (or, equivalently, the shading factor) at every observed bid.

3.2.2 Estimating the Value of Customer Order Flow

The data set from Canadian Treasury auctions containing all submitted bids is analyzed in Hortaçsu and Kastl (2012). The link between the observed bids and the unobserved willingness-to-pay, is given by equation (2). As mentioned before, a unique feature of the Canadian data set is that it contains timestamps of individual bids, and equally importantly, it also contains all bids that are before the auction deadline superseded by updated bids - and this is the detail of the auction that I want to focus on. In particular, primary dealers frequently submit bids that they change at a later point, often after they have been asked to submit bids on a customer's behalf. Hortaçsu and Kastl (2012) provide a simple model rationalizing the different timing of observed bid submissions. In its simplest form, this requires introducing uncertainty in whether customers' bids arrive at all and if they do, whether the subsequently updated dealer bid will make it in time to make the deadline.

Figure 1 in Hortaçsu and Kastl (2012) shows that many customers wait literally until the very last minute before submitting their bids. Such "last minute" bidding behavior by customers can be rationalized as a strategic response by customers who do not want dealers to utilize the information in their bids. There may be reasons for customers to voluntarily share their information with dealers as well. For example, Bloomberg Business published an article on 4/4/2013 where a representative of Blackrock, the world's largest asset manager, described why Blackrock chooses to participate indirectly by submitting bids through PDs: *"While we can go direct, most of the time we don't. We feel that the dealers provide us with a lot of services. Our philosophy at this point is, to the extent we can share some of that information with trusted partners who won't misuse that information, we prefer to reward the primary dealers that provide us all that value."*[1] In previous work studying the Canadian Treasury market, Hortaçsu and Sareen (2006) find that some dealers' modifications to their own bids in response to these late customer bids narrowly missed the bid submission deadline, and that such missed bid modification opportunities had a negative impact on dealers' ex-post profits.

With slight modifications on top of our auction model from Section 2.1 that allow for customers and primary dealers submitting bids in various times, we can use equation (2) and the method described in Section 2.2 to estimate willingness-to-pay that would rationalize each observed bid. With these estimates in hand, it is natural to proceed by comparing expected profits corresponding to a primary dealer's bid that was submitted absent any customer's order information to expected profits corresponding to the bid that was submitted after observing customers' order flow (both within the same auction). While intuitive, this approach would be premature. In Assumption 3, we have imposed that bidders' values are private – in particular, we assumed that if a primary dealer were to learn a customer's signal by observing that customer's bid, she would not update her estimate of the value. This assumption is often questioned and Treasury Bills are often mentioned as an example of objects that should be modeled using a model with interdependent values. Given

[1] `http://www.bloomberg.com/news/articles/2013-04-04/bond-traders-club-loses-cachet-in-most-important-market`.

the data available in the Canadian Treasury auctions, we can formally test this assumption, at least as far interdependency of values between primary dealers and their customers is concerned.

3.2.3 Private versus Interdependent Values in Treasury Bill Auctions

Recall that in the Canadian Treasury Bill auctions the researcher observes instances where a bidder (primary dealer) submits two bids in the same auction, one before customers' order flow arrives and one after it arrives. This setting thus provides us with a natural laboratory where to investigate whether primary dealers are just learning about competition in the upcoming auction (and not updating their valuation estimates) or whether the PDs may also be learning about fundamentals, and hence update their value estimates after observing customers' bids. When deciding how to translate willingness-to-pay into the bid, the primary dealer must form beliefs about the distribution of the market clearing price (as Equation (2) reveals). To form such beliefs, the primary dealer integrates over all uncertainty: the signals of rival primary dealers, the signals of all customers, which primary dealer a customer might route her bid through etc. By observing a customer's order, part of this uncertainty gets resolved: a primary dealer can therefore update her belief about the distribution of the market clearing price by evaluating (6). Since this updating can be easily replicated in the estimation by appropriately adapting the resampling technique, it allows us to form a formal hypothesis test about whether values are indeed private. In particular, let $\left\{v_k^{BI}\left(q_k,\theta\right)\right\}_{k=1}^{K}$ denote the vector of the estimated willingness-to-pay that rationalizes the observed vector of bids before the customer's order arrives (hence, "BI"), $\left\{b_k^{BI}\left(q_k,\theta\right)\right\}_{k=1}^{K}$, and $\left\{v_k^{AI}\left(q_k,\theta\right)\right\}_{k=1}^{K^{BI}}$ denote the vector of the estimated willingness-to-pay that rationalizes the observed vector of bids after the customer's bid arrives (hence, "AI"), $\left\{b_k^{AI}\left(q_k,\theta\right)\right\}_{k=1}^{K^{AI}}$. If we were to observe a bid for the same quantity being part of both the bid before customer's information arrive and the bid after its arrival, we can simply formulate a statistical test at that quantity. Let $T_j\left(q\right) = v^{BI}\left(q,\theta\right) - v^{AI}\left(q,\theta\right)$ be the difference between the rationalizing willingness-to-pay for a given quantity, q, before and after information arrives (taking into account the updating about the distribution of the market clearing price during the estimation). Testing the null hypothesis (of no learning about fundamentals) then involves testing that $T_j(q) = 0, \forall j, q$. Of course, the null hypothesis is essentially a composite hypothesis of all modeling assumptions (e.g., in addition to private values, we also assume independence etc.). To test the hypotheses jointly, since there is no result on the most powerful test, we can define several joint hypotheses tests and perform them concurrently. For example, the following test is motivated by the well-known χ^2-test, which can be additionally standardized by the standard deviation of the asymptotic distribution of each individual test statistic.

$$T = \sum_{(i,t)=(1,1)}^{(N,T)} T_{i,t}\left(q\right)^2 \tag{8}$$

Hortaçsu and Kastl (2012) report that the null hypothesis of no learning about fundamentals from customer's bids can neither be rejected based on (8), nor based on several alternative tests. While this still does not preclude interdependency of values between primary dealers themselves, given that many of the customers are large players (such as Blackrock), this evidence is at least suggestive that modeling the information structure in Treasury auctions as private values is reasonable. It is important to note that the crucial point about the information structure is not that the security of interest might have common unknown value at the secondary market after the auction, but rather whether or not the primary dealers have different information about this ex-post value before the auction. If the information is symmetric, and potentially imperfect, the heterogeneity in values (and thus in bids) might still be attributable purely to heterogeneity in the private component of the values.

3.2.4 Putting It Together: Quantifying the Order Flow

Since I argued in the previous subsection that the private value assumption might be reasonable in the context of our application, let me start quantifying the value of customers' order flow information. We will make use of the dynamics of bidding within an auction. Since the bidding takes place typically within few minutes before the auction's deadline, I will assume that nothing else is changing other than dealer observing the customers' orders. This allows me to quantify the value of order flow by calculating the expected surplus (profits) of a primary dealer associated with the original bid (before customers' order information arrives) and contrast it against the expected profits associated with the updated bid which includes the information about the customer's bid.

Let $\Pi^I(s_i, z_i)$ denote the expected profit of a dealer d, when using the bidding strategy $y^I(p, s_i, z_i)$, i.e., after incorporating the information from customers' orders which is contained in the realization of a random variable Z_i. Similarly, let $\Pi^U(s_i,)$ denote the expected profit corresponding to the bidding strategy $y^U(p, s_i)$, i.e., before customers' orders arrive.

The value of information in terms of this notation is as follows:

$$VI = \int_0^\infty \Pi^I(s_i, z_i)\, dH\left(P^c, y^I(s_i, z_i)\right) - \int_0^\infty \Pi^U(s_i)\, dH\left(P^c, y^U(s_i,)\right) \tag{9}$$

where $H(P^c, y^x(\cdot))$ is the distribution of the market clearing price, P^c, given other bidders using equilibrium strategies and dealer d following the strategy $y^x(\cdot)$, where the superscript $x \in \{I, U\}$ indexes the "informed" and the "uninformed" dealer. Recall that the distribution of the market clearing price is different for uninformed and informed dealer: the former needs to integrate out the uncertainty with respect to all customers, whereas the latter takes one customer's realized bid as given and only integrates out the remaining uncertainty.

Hortaçsu and Kastl (2012) report that about one third of PDs' profits can be attributed to the customers' order flow information. Overall, however, the auctions seem fairly competitive and the profits thus do not seem excessive. Hortaçsu et al. (2018) reached a similar conclusion for the US

Treasury auctions.

Although the calculation above allows us to measure how much extra (expected) profit dealers make when updating their bids, we do not attempt to answer the following, more ambitious question: what is the value to dealers from the institutional structure that allows them to observe and react to customer bids? Answering this question would require us to calculate dealer profits under the counterfactual scenario where dealers do not observe customer bids, for which we would have to recompute the equilibrium bidding strategies of customers and dealers. Unfortunately, computing equilibrium strategies in (asymmetric) discriminatory multi-unit auctions is still an open question, and we will leave this calculation to future research.

The main takeaway from the Canadian Treasury auctions is thus that (i) customers' order flow is an important source of primary dealers' rents (up to one third in the Canadian auctions) and (ii) private values may be a reasonable approximation of the information structure in Treasury bill auctions.

Another important question is why indirect bidders voluntarily choose to submit their bids through a primary dealer rather than to participate in the Treasury auctions directly. Quantifying the benefits and costs of the primary dealer system formally would be a very valuable and ambitious research project. Since the primary dealers likely provide many complementary services to their customers, ignoring these channels would lead to an inconsistent estimate of rents PDs and customers accrue just from the Treasury auctions. For example, PDs might step in when the customers need to off-load some illiquid asset, they might have valuable information about demand for such assets and thus may act as intermediaries. In other markets PDs may engage in trade in order to gather information about the fundamentals of assets for which for reasons of illiquidity or for other reasons prices may not aggregate information properly (as in Brancaccio et al. (2018)) etc. While being a primary dealer also has its costs in terms of duties (such as the commitment to regularly and actively participate in Treasury auctions) or strict regulatory requirements, it entails also many important benefits. In addition to the benefits of observing customer order flow discussed above, primary dealers typically get exclusive access to important and advantageous sources of liquidity, such as quantitative easing operations etc.

3.3 From Auctions to a Demand System for Government Securities

Demand estimation is part of the essential toolbox of any IO economist, and it has been one of the central pieces of research agenda over the past four decades (Deaton and Muellbauer (1980); Berry et al. (1995)). It is easy to see why: whenever we are interested in counterfactuals, for which observed data are insufficient on their own, we need to complement the data with an economic model. For example, we want to evaluate a policy that would amount to subsidizing one of the goods. A model of such counterfactual world would typically involve both a demand and a supply side - and depending on the counterfactuals of interest we may need one, the other, or both. For many questions of interest (such as a subsidy analysis, a merger analysis or valuation of new products), knowledge of the matrix of own- and cross-price elasticities and the substitution matrix are necessary, so that we can calculate equilibria in the counterfactual worlds given some assumptions on the nature of the competition. With a lot of products, the demand system in the product space, where demand for each product depends on quantities purchased of all other products, involves a lot of parameters (proportional to the square of the number of products). Moreover, estimation requires a lot of instruments to cope with the classic endogeneity problem: since firms choose prices optimally, they tend to pick higher prices in response to high (unobserved) demand shocks. One important simplification (Lancaster (1966)) is to map products onto a space of characteristics and consumers making choices between bundles of characteristics. Berry et al. (1995) propose a variant of this demand estimation approach for aggregate data, and show how to recover preferences from aggregate market shares of individual goods. Koijen and Yogo (2019) adapt this approach to financial markets. It is important to note that this literature mostly deals with estimating price elasticities for goods that are substitutes. While incorporating complementarities is feasible by introducing additional parameters as in Gentzkow (2007), the main benefit of the characteristic space approach, namely the decrease of the number of parameters, is diminished.

In a typical demand estimation setting a researcher observes one point on the demand curve at any point in time, i.e., the equilibrium price and quantity. Hence one typically needs to utilize exogenous variation in prices that is attributable to shifts in the supply curve (such as shocks to the marginal costs of production) in order to link the observed points together to be able to estimate the demand curve - this essentially is the standard endogeneity problem mentioned above. In our setting, we observe several points on the demand curve at any point in time: one point corresponds to every bid. This substantially simplifies the problem relative to the applications based on the tools developed in Berry et al. (1995) or Koijen and Yogo (2019).

Why are we interested in knowing the substitution patterns between securities of different maturities? In case of Treasury bill auctions, the offices in charge of debt management (such as the Office of Debt Management in the U.S.) need to regularly come up with an issuance plan: how much of each maturity to issue in order to raise some target revenue and also in order to keep financial markets happy (e.g., by providing sufficient amounts of the most sought after maturities) and steer the yield curve if so desired. Since demands are interdependent, increasing supply of one maturity and decreasing supply of another is not revenue neutral as it impacts the expected allocations of both. Therefore, this choice also impacts the (expected) marginal willingness-to-pay - and, indirectly, the bids in the auctions. Hence, it can both increase and decrease revenue and we would need an estimate of the full system of preferences of primary dealers over these securities. This is the goal of the application we will discuss in this section.

Yet another important source of interdependency of demands for various securities arises from there being auctions in different countries and essentially the same (or substantially overlapping) set of dealers participates in all of them. This direction is left for future research, but in the last section of this paper I briefly discuss an application which illustrates using the QE auctions that such cross-border dependencies are indeed important.

To achieve this goal, I will make use of another interesting feature of Canadian Treasury bill auctions (and in fact, of many primary markets around the world) analyzed in Allen et al. (2018). The auctioneer typically conducts several auctions simultaneously. In case of Canada, there are auctions for 3-month, 6-month and 12-month Treasury bills that are run at the same time. Since typically all bidders (especially the primary dealers) participate in all three auctions, it is not difficult to see that the auctions will be interdependent. The willingness-to-pay for maturity m will depend on the expectations of how much of the other maturities a dealer might (expect to) win. Since through the bid inversion technique discussed in previous section we can recover the willingness-to-pay rationalizing each bid, we can for any triplet of auctions recover a corresponding triplet of willingnesses-to-pay and hence we can study how these depend on each other.

Allen et al. (2018) generalize the methods for estimating demands from bidding data to allow for interdependencies along these lines, and find that 3, 6 and 12-month bills are often complementary in the primary market for Treasury bills. They also present a model that captures the interplay between the primary and secondary markets to provide a rationale for these seemingly counter-intuitive findings.

To be more precise, let a simultaneous auction be indexed by τ and let bidder i of type $s^g_{m,i,\tau}$ in bidder group $g \in \{d = \text{dealer}, c = \text{customer}\}$ have the following willingness to pay for amount q_m in auction m conditional on winning q_{-m} of the other two maturities and keeping a share $(1 - \kappa_m)$ on its own balance sheet

$$v_m(q_m, q_{-m}, s^g_{m,i,\tau}) = \alpha + (1 - \kappa_m)s^g_{m,i,\tau} + \lambda_m q_m + \delta_m \cdot q_{-m}. \tag{10}$$

δ_m measures the interdependencies across maturities. For example, If $\delta_{3M,6M} < 0$, bidders are willing to pay less for any amount of the 3-month maturity the more they purchase of the 6-month bills, hence the bills are substitutes. When $\delta_{3M,6M} > 0$ they are complementary, and independent if $\delta_{3M,6M} = 0$. This functional form is as flexible in terms of the allowed substitution patterns as the approach by Gentzkow (2007).

Given the rules of the auction, bidders cannot express their bids for any given maturity as conditional on quantities of other maturities won - they are only allowed to submit bids which only condition the offered payment on the quantity of that particular maturity. In such an environment, even the bid under truthful bidding, let me denote it by $\tilde{v}_m(q_m, s^g_{m,i,\tau})$, would not correspond to the true MWTP $v_m(q_m, q_{-m}, s^g_{m,i,\tau})$. This is because a dealer's actual marginal benefit from winning amount q_m depends on how much he will win of the other assets: $q^*_{-m,i}$. Since auctions take place in parallel, these random quantities $q^*_{-m,i}$ need to be integrated out:

$$\tilde{v}_m(q_m, s^g_{m,i,\tau}) = \mathbb{E}[v_m(q_m, q^*_{-m,i}, s^g_{m,i,\tau})| \text{ win } q_m].$$

Hence, following our bid inversion technique, after the first stage of the estimation procedure we recover $\tilde{v}_m(q_m, s^g_{m,i,\tau})$. In addition, we estimate the joint distribution of market clearing prices (and quantities) which allows us to estimate the conditional expectation $\mathbb{E}[q^*_{-m,i}| \text{ win } q_m]$. Given the linearity assumption in (10) we can then estimate the parameters of interest, δ_m, by a simple projection: a linear regression with bidder-auction-time fixed effects that control for $\alpha + (1 - \kappa_m)s^g_{m,i,\tau}$.

Allen et al. (2018) find non-negligible interdependencies in the primary dealers' demands for government securities. There is quite a bit of heterogeneity among dealers, but roughly there are two main "types." For one group of dealers, treasuries of different maturities are complementary to each other. For these dealers, money market is roughly their primary business. They have lots of clients and run large money desks. For the second group of dealers, the treasuries of different maturities are substitutes to each other. These dealers have comparatively smaller money desks and the focus of their business models is on other financial markets.

Aggregating these individual demands into a market demand results in significant dependencies, although not as clear cut as at the individual dealer level. Allen et al. (2018) estimate that the marginal valuation for a 3-month T-bill increases when going from an allocation with no other T-bills of other maturities at all to an allocation corresponding to the average observed allocations of 6 and 12-month bills (about 200 million each) by about 0.5 basis points. While this effect may seem small in absolute size this "cross" effect is actually relatively large compared to the "own" effect: the marginal value for the 3-month bill drops by about 1 basis point when going from none to an allocation of 400 million.

4 Auctions and Quantitative Easing

In the aftermath of the financial crisis of 2008 many central banks engaged in what became known as "unconventional monetary policy." When interest rates drop to zero, to avoid the liquidity trap central banks engaged in open market operations to expand the money supply. Some central banks were buying privately issued securities (such as mortgage-backed securities) and some instead focused on buying back long term (or other rather illiquid) government bonds. This buyback was often implemented via an auction (Han et al. (2007)). This was also the case in the United Kingdom since March 2009 (see Joyce et al. (2011) for details). The purchases amounted to over £400 billion. For the case of the US Song and Zhu (2014) analyze data on winning bids from the US QE auctions that were made publicly available. The fact that losing bids are not observed, however, does not allow the authors to apply the methodology described in this paper since the losing bids are necessary to estimate bidders' beliefs (under rational expectations) about how price and probability of winning a certain allocation were to change with a perturbation of the bid. Boneva et al. (2018) analyze data from the UK that include also the losing bids and therefore a structural analysis along the lines described in the earlier sections of this paper can be applied.

4.1 QE Auctions in the UK

The Bank of England's version of quantitative easing used a multi-object, multi-unit, discriminatory price reverse auction format and hence the machinery described in the previous sections is once again a good tool to apply. Separate auctions were run for different maturity sectors and unlike in the Canadian Treasury bill auctions discussed above these auctions took place on different days. Initially, the Bank ran only two auctions per week, one for maturities between 5 and 10 years, and one for maturities between 10 to 25 years. This was later expanded to include even longer-term gilts. A list of eligible bonds was published ahead of the auction.

At the close of the auction at 2:45pm, the received offers were ranked according to the spread between the offered yield and the secondary-market yield prevailing at 2:45pm as observed by the Bank of England on the Reuters screen. The auctioneer then accepted best offers up to the pre-announced volume. Successful auction participants received their asking price in exchange for the securities, meaning that all purchases are undertaken on a discriminatory price basis. During each auction, the Bank monitored the developments in the secondary market for the eligible bonds for unusual price developments as the auction format of course might tempt some participants to manipulate the price of certain securities in order to take advantage of the QE auction.

To map this mechanism into our multi-unit share auction framework, Boneva et al. (2018) assume that bidders have rational expectations and hence (on average) accurately forecast the secondary market prices prevailing at the auction close. With this assumption, one can normalize the submitted offers by this price and hence obtain "homogenized" offers that are comparable across auctions and across eligible bonds. One can then follow the method described above to assign an estimated marginal value (or marginal willingness-to-accept in the case of these reverse auctions, where bidders are making offers to sell the securities to the auctioneer), that rationalizes the observed offer.

4.2 Results from the UK QE Auctions

Having recovered the marginal willingness-to-accept for every offer in the QE auctions, Boneva et al. (2018) study the determinants of these marginal values by projecting them on many covariates.They find that inventory risk (duration of the various bonds in their portfolios) and regulatory constraints affect the primary dealers bidding and are indeed among the most important determinants of the marginal willingness-to-accept.

A more important and perhaps also more surprising result involves the spillovers of unconventional monetary policy effects across borders. Since many primary dealers operate in most of the important world markets, Boneva et al. (2018) also investigate whether QE implemented in the US spills over to the UK - in terms of affecting primary dealers' willingness-to-accept through channels such as a global budget and they find that it indeed does.

5 Conclusion

The focus of this article has been on pointing out how seemingly innocuous details of auction mechanisms might help us gain insights about the working of the markets that are otherwise hard to analyze. The first key step is to build a theoretical model that closely follows the rules of the actual market place. This does not only involve modeling the payment and allocation rules (which are of course the two essential ingredients of any mechanism), but also the information structure underlying the problem at hand. The solution of the appropriately written down model allows us then to construct a link between the observables and unobservables and hence to recover the unobserved objects of interest, such as marginal values. The applications considered in this paper included estimating the value of customer order flow, which made use of (observed) dynamics of bids before the auction deadline. Primary dealers adjust their bids in response to the arrival of information (bids of their customers). I also illustrated how to estimate a demand system for financial securities by primary dealers. Finally, I discussed how to potentially evaluate the impact of non-conventional monetary policy on secondary markets and quantify the externalities such policies might create across borders.

References

Allen, Jason, Jakub Kastl, and Milena Wittwer, "Identifying Dependencies in the Demand for Government Securities," September 2018. working paper.

Athey, Susan and Philip A. Haile, "Identification of Standard Auction Models," *Econometrica*, 2002, *70* (6), pp. 2107–2140.

Ausubel, Lawrence, Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka, "Demand Reduction and Inefficiency in Multi-Unit Auctions," *Review of Economic Studies*, 2014, *81* (4), pp. 1366–1400.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *63* (4), pp.841–890.

Boneva, Lena, Jakub Kastl, and Filip Zikes, "Dealer balance sheets and bidding behavior in the UK QE reverse auctions," 2018. working paper.

Brancaccio, Giulia, Dan Li, and Norman Schürhoff, "Learning by Trading: The case of the US Market for Municipal Bonds," 2018. working paper.

Bresnahan, Timothy F., "Empirical Studies of Industries with Market Power," in R. Schmalensee and R.D. Willig, eds., *Handbook of Industrial Organization*, Vol. 2, Elsevier, 1989.

Cassola, Nuno, Ali Hortaçsu, and Jakub Kastl, "The 2007 Subprime Market Crisis in the EURO Area Through the Lens of ECB Repo Auctions," *Econometrica*, 2013, *81* (4), pp. 1309–1345.

Deaton, Angus and John Muellbauer, "An Almost Ideal Demand System," *The American Economic Review*, June 1980, *70* (3), pp. 312–326.

Gentzkow, Matthew, "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 2007, *97* (3), pp. 713–744.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 1985, *14* (1), pp. 71–100.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong, "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 2000, *68* (3), pp. 525–574.

Han, Bing, Francis Longstaff, and Craig Merrill, "The U.S. Treasury buyback auctions: the cost of retiring illiquid bonds," *The Journal of Finance*, 2007, *62* (6), 2673–2693.

Hendricks, Kenneth and Robert H. Porter, "An Empirical Study of an Auction with Asymmetric Information," *The American Economic Review*, 1988, *78* (5), pp. 865–883.

Hortaçsu, Ali, "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," 2002. working paper.

_ **and David McAdams**, "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," *Journal of Political Economy*, 2010, *118* (5), pp. 833–865.

_ **and Jakub Kastl**, "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions," *Econometrica*, 2012, *80* (6), pp.2511–2542.

_ **and Samita Sareen**, "Order Flow and the Formation of Dealer Bids in Treasury Auctions," 2006. working paper.

_ **, Jakub Kastl, and Allen Zhang**, "Bid Shading and Bidder Surplus in U.S. Treasury Auctions," *The American Economic Review*, 2018, *108* (1), pp. 1–24.

Jovanovic, Boyan, "Observable Implications of Models with Multiple Equilibria," *Econometrica*, 1989, *57* (6), 1431–1437.

Joyce, Michael, Matthew Tong, and Robert Woods, "The United Kingdoms quantitative easing policy: design, operation and impact," *Bank of England Quarterly Bulletin*, 2011, *51* (3), 200–212.

Kastl, Jakub, "Discrete Bids and Empirical Inference in Divisible Good Auctions," *Review of Economic Studies*, 2011, *78*, pp. 978–1014.

_ , "On the Properties of Equilibria in Private Value Divisible Good Auctions with Constrained Bidding," *Journal of Mathematical Economics*, 2012, *48* (6), pp. 339–352.

_ , *Advances in Economics and Econometrics: Eleventh World Congress*, Cambridge University Press, 2017.

Koijen, Ralph S.J. and Motohiro Yogo, "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, 2019, *127* (4), pp. 1475–1515.

Krasnokutskaya, Elena, "Identification and Estimation of Auction Models with Unobserved Heterogeneity," *The Review of Economic Studies*, 2011, *78* (1), 293–327.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

_ , "Informed Speculation with Imperfect Competition," *The Review of Economic Studies*, 1989, *56* (3), pp. 317–355.

Laffont, Jean-Jacques and Quang Vuong, "Structural Analysis of Auction Data," *The American Economic Review*, 1996, *86* (2), pp. 414–420.

Lancaster, Kelvin J., "A New Approach to Consumer Theory," *Journal of Political Economy*, 1966, *74* (2), pp.132–157.

McAdams, David, "Partial Identification and Testable Restrictions in Multi-Unit Auctions," *Journal of Econometrics*, 2008, *146*, pp. 74–85.

Milgrom, Paul R. and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, *50* (5), pp. 1089–1122.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 1989, *57* (5), pp.1027–1057.

Pycia, Marek and Kyle Woodward, "Pay-As-Bid: Selling Divisible Goods to Uninformed Bidders," January 2015. mimeo.

Rostek, Marzena and Marek Weretka, "Price Inference in Small Markets," *Econometrica*, 2012, *80* (2), pp.687–711.

Shleifer, Andrei, "Do Demand Curves for Stocks Slope Down?," *The Journal of Finance*, July 1986, *41* (3), pp.579–590.

Song, Zhaogang and Haoxiang Zhu, "QE Auctions of Treasury Bonds," *Finance and Economics Discussion Series, Federal Reserve Board*, 2014.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 1961, *16*, pp. 8–37.

Vives, Xavier, "Strategic Supply Function Competition with Private Information," *Econometrica*, 2011, *79* (6), pp.1919–1966.

Wilson, Robert, "Auctions of Shares," *The Quarterly Journal of Economics*, 1979, *93* (4), pp. 675–689.