

NBER WORKING PAPER SERIES

BID SHADING AND BIDDER SURPLUS IN THE U.S. TREASURY AUCTION SYSTEM

Ali Hortaçsu  
Jakub Kastl  
Allen Zhang

Working Paper 24024  
<http://www.nber.org/papers/w24024>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
November 2017

The views expressed in this paper are those of the authors and should not be interpreted as reflecting the views of the U.S. Department of the Treasury or the National Bureau of Economic Research. We thank Philip Haile, Kenneth Hendricks, Darrell Duffie, Paulo Somaini, Terry Belton and participants of the U.S. Treasury Roundtable, 2015 NBER IO Summer Meetings and 2015 NBER Market Design Meeting for their valuable comments. Kastl is grateful for the financial support of the NSF (SES-1352305) and the Sloan Foundation, Hortaçsu the financial support of the NSF (SES-1124073, ICES-1216083, and SES-1426823). All remaining errors are ours.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24024.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Ali Hortaçsu, Jakub Kastl, and Allen Zhang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Bid Shading and Bidder Surplus in the U.S. Treasury Auction System  
Ali Hortaçsu, Jakub Kastl, and Allen Zhang  
NBER Working Paper No. 24024  
November 2017  
JEL No. G12,L1

**ABSTRACT**

We analyze bidding data from uniform price auctions of U.S. Treasury bills and notes between July 2009-October 2013. Primary dealers consistently bid higher yields compared to direct and indirect bidders. We estimate a structural model of bidding that takes into account informational asymmetries introduced by the bidding system employed by the U.S. Treasury. While primary dealers' estimated willingness-to-pay is higher than direct and indirect bidders', their ability to bid-shade is even higher, leading to higher yield/lower price bids. Total bidder surplus averaged to about 3 basis points across the sample period along with efficiency losses around 2 basis points.

Ali Hortaçsu  
Department of Economics  
University of Chicago  
1126 East 59th Street  
Chicago, IL 60637  
and NBER  
hortacsu@uchicago.edu

Allen Zhang  
Department of the Treasury  
xpzhang2001@gmail.com

Jakub Kastl  
Department of Economics  
Princeton University  
Princeton, NJ 08544-1021  
and NBER  
jkastl@princeton.edu

# 1 Introduction

In 2013, the U.S. Treasury auctioned 7.9 trillion dollars of government debt to a global set of institutions and investors (TreasuryDirect.gov). The debt was issued in a variety of instruments, covering bills (up to 1 year maturity), 2-10 year notes, 30 year coupon bonds, and Treasury Inflation-Protected Securities (TIPS). The mandate of the U.S. Treasury is to achieve the lowest cost of financing over time, taking into account considerable uncertainty in the borrowing needs of the government and demand for U.S. Treasuries by investors.<sup>1</sup> Treasury also seeks to “facilitate regular and predictable issuance” across a range of maturity classes. To this end, the Treasury adopted auctions as their preferred method of marketing short term securities in 1920 (Garbade 2008), and auctions became the preferred method of selling long-term securities in the 1970s (Garbade 2004). The Treasury employed a discriminatory/pay-as-bid format until 1998, when, after some experimentation with its 2- and 5-year note auctions in 1992, the uniform price format was adopted as the method of sale (Malvey and Archibald 1998).

A large theoretical literature, dating back to the classic work of Friedman (1960) and Vickrey (1961), studies the design of Treasury auctions. The design of the auction mechanism becomes an important question especially when individual bidders possess market power, as the exercise of market power leads to deviations from efficient allocations (Ausubel, Cramton, Pycia, Rostek and Weretka (2014)). The competitiveness of the Treasury auction process and the design of the auction system has come under the spotlight a number of times, including the 1991 Salomon Brothers “short squeeze” scandal (see e.g. Jegadeesh (1993)), and the discussions that led to Treasury’s switch to the uniform price auction mechanism in November 1998 (Malvey and Archibald 1998).

While there is a large theoretical literature studying strategic bidder behavior, there has been very little empirical work on the U.S. Treasury auction system that studies the behavior of individual bidders. This paper hopes to bridge this empirical gap by studying individual bid level data. Specifically, our aim is to adapt models of strategic behavior to the particular institutional setting of the U.S. Treasury auction market, and to develop quantitative benchmarks for assessing

---

<sup>1</sup>Peter Fisher, then Under Secretary of Treasury for Domestic Finance, in a speech titled “Remarks before the Bond Market Association Legal and Compliance Conference” on January 8, 2002, stated that “the overarching objective for the management of the Treasury’s marketable debt is to achieve the lowest borrowing cost, over time, for the federal government’s financing needs.”

the competitiveness and cost-effectiveness of this important marketplace.

We start, in Section 2, by describing differences in bid patterns across the three different classes of bidders: primary dealers, direct bidders (financial institutions who are not primary dealers but submit bids directly into the auction), and indirect bidders (financial institutions who submit bids through primary dealers). As detailed in this section, we find that primary dealers systematically bid lower prices (higher yields). What can drive this systematic difference? It is possible that primary dealers have lower valuations for the Treasury securities, and this lower valuation translates into lower bids. However, it is also possible that primary dealers possess market power, and are able to exercise this market power by submitting (price) bids that are below their true valuations.

We thus provide, in Section 3, a model-based decomposition of the observed bids into the bidder's true willingness-to-pay and a "bid shading"/market power component. Our theoretical model adapts the seminal "share auction" model of Wilson (1979), in which bidders submit demand schedules as their bids. This model captures the strategic complexity of the Treasury's uniform price auction mechanism very well, and it is in many ways related to classic models of imperfect competition such as Cournot. In particular, consider a setting (depicted in Figure 1) where an oligopsonistic bidder with downward sloping demand for the security knows the residual supply function that she is facing, and is allowed to submit a single price-quantity point as her bid. Following basic monopsony theory, this bidder will not select the competitive outcome  $(P^{comp}, Q^{comp})$ , which is the intersection of her demand curve and the residual supply curve. Instead, she has the incentive to "shade" her bid, and pick a lower price-quantity point on the residual supply curve, such as  $(P^*, Q^*)$ . This gives her higher surplus (the gray shaded area below her demand curve up to  $(P^*, Q^*)$ ) than if she had bid the competitive price and quantity.

The ability of this bidder to "shade" her bid depends on the elasticity of the residual supply curve she is facing. If this were a small bidder among many others, the residual supply she would be facing would essentially be flat, allowing for very little ability to bid-shade, as decreasing the quantity demanded would not result in any appreciable change in the market clearing price. The optimal bid then is to bid one's true demand curve. Therefore, the key identifying variation in the data is the (expected) shape of the residual supply curves that a bidder may face around the

point where they intersect this bidder’s bid. This, in turn, is driven by shapes of other bidders’ bids, which are observed in the data without resorting to behavioral assumptions. For example, if all bids were identical (and the auction were oversubscribed), no bidder would be able to affect the market clearing price individually and hence everyone would be bidding their values. The more heterogeneity in submitted bids and the lower the elasticity of the expected residual supply, the larger the bid shading/market power component. Note, of course, that quantifying the exact magnitude of bid shading exercised by each bidder depends on the assumption of unilateral and noncooperative expected profit maximization.

The Wilson (1979) model, and its generalization that we discuss here allows bidders to have private information about their true demands/valuations for the securities and to submit more than one price-quantity pair as their bids. This induces uncertainty in competing bids, and thus an uncertain residual supply curve. What Wilson derives is a locus of price-quantity points comprising a “bid function” that maximizes the bidder’s expected surplus against possible realizations of the residual supply curve.

A limitation of Wilson’s model is that bidders are assumed to bid continuous bid functions. However, in most real world settings, bidders are confined to a discrete strategy space that limits the number of price-quantity points they can bid. Indeed, in the auctions that we study, bidders utilize, on average, 3 to 5 price-quantity points, with a nontrivial fraction of bidders submitting a single step, making “discreteness” a particularly important problem. To deal with this issue, we adopt the generalization of the Wilson model by Kastl (2011).

Yet another complication we face in this setting is the fact that bidders have inherently asymmetric information sets. This asymmetry is introduced by the fact that some bidders (the “primary dealers”) route the bids of others (the “indirect” bidders), and hence observe a part of the residual supply curve that other bidders do not. Following our prior work on Canadian treasury auctions (Hortaçsu and Kastl (2012)) which had a very similar feature, our model also incorporates the reduced uncertainty/more precise information that primary dealers possess regarding residual supply.

We thus characterize bidders’ optimal decisions in terms of their (unobserved) “true” demands/marginal valuations, and their beliefs about the distribution and shape of residual sup-

ply. Indeed, the optimality condition closely resembles the inverse elasticity markup rule of classical monopoly theory. The optimality condition thus allows us to infer the unobserved demands/marginal valuations of the bidders, provided that we make certain assumptions regarding bidders’ beliefs. We follow, as in some of our previous work (Hortaçsu (2002), Hortaçsu and McAdams (2010), Kastl (2011), and Hortaçsu and Kastl (2012)) the seminal insight of Guerre, Perrigne and Vuong (2000)<sup>2</sup> that under the assumption that bidders are playing the Bayesian-Nash equilibrium of the game, one can use the realized distribution of residual supplies as an empirical estimate of bidders’ ex-ante beliefs. Once we have the bidders’ beliefs, we can invert the optimality condition to recover bidders’ unobserved demand/marginal valuations rationalizing their observed behavior.

In Section 5, we estimate individual bidders’ “rationalizing” marginal valuations/demand curves to answer two sets of questions. The first is to quantify the extent of market power exercised by bidders through bid-shading, as in Figure 1, and how this varies across different subclasses of bidders. We find that primary dealers shade their bids more than direct and indirect bidders – which goes towards explaining the fact that primary dealers bid significantly lower prices/higher yields than direct and indirect bidders.

The second question we answer is regarding bidder surplus. In Figure 1, one can quantify the bidder surplus by looking at the area under the bidder’s true demand curve above  $P^*$  and up to  $Q^*$ . Since we have estimates of bidders’ “rationalizing” demand curves, we can calculate this area for each bidder in our data. Our main finding here is that primary dealers, perhaps not surprisingly, extract more surplus from the auctions than direct or indirect bidders. The magnitude of the surpluses vary quite a bit between maturity classes, with very little bidder surplus in Treasury bills, and larger surpluses for Treasury notes.

Our surplus calculations also allow us to connect to the long literature studying the “optimal auction mechanism” to use to sell Treasury securities. This literature dates at least back Friedman (1960), who pointed out the bid-shading incentives of bidders in the then-used discriminatory/pay-as-bid auction utilized by the Treasury, and advocated the use of the uniform-price auction, as it would alleviate the bid-shading incentive of smaller bidders (for whom bidding one’s valuation is

---

<sup>2</sup>And also of Elyakime, Laffont, Loisel and Vuong (1997) and Laffont and Vuong (1996).

approximately optimal) and lower their cost of participating in these auctions. However, as noted by several authors including Wilson (1979) and Ausubel et al. (2014), a bid-shading incentive remains for larger bidders in the uniform price auction. Ausubel et al. (2014) also show that optimal bid shading in these auctions distorts the efficiency of the allocations as well, and thus a general ranking of expected revenues from discriminatory and uniform price auctions can not be made without knowledge about the specific features of bidders' demands.

Given the theoretical vacuum, a variety of empirical approaches have been employed to assess the efficacy of Treasury auction mechanisms. The Treasury's own study of this question, as reported by Malvey and Archibald (1998), was based on experimentation with the uniform price format for 2- and 5-year notes. To assess the revenue properties of the uniform vs. the status-quo discriminatory format, Malvey and Archibald calculated the auction-when-issued rate spread, and did not statistically reject a mean difference across the different auction formats. However, they note that the uniform price auctions "produce a broader distribution of auction awards" across bidders, and especially a lowered concentration of awards to top primary dealers.

Our empirical approach differs from that of Malvey and Archibald's and related studies, in that we do not look at when-issued or secondary market prices to assess the "value" of the securities being sold.<sup>3</sup> Indeed, what we are interested is the "inframarginal surplus" of the bidders, which, in the presence of downward sloping demand, will not be apparent from looking at market clearing prices either in the primary or secondary markets. Heterogeneity in valuations that lead to downward sloping demand for these securities may arise from many different sources: buy-and-hold bidders may have idiosyncratic portfolio immunization needs, financial intermediaries may attach different valuations to the Treasuries due to e.g. their use as collateral, primary dealers may value having an inventory of Treasury beyond its resale value because being a primary creates additional value

---

<sup>3</sup>We looked at the differential between auction-close when-issued rates and the auction stop-out rate using the best bid-side quote on TradeWeb 15 minutes before the auction close. We found that for the majority of maturities and years considered, our sample of auctions were "through," i.e. the when-issued yield (price) was systematically higher (lower) than the stop-out yield (price). This may be consistent with the explanation given by Bikhchandani, Edsparr and Huang (2000) that primary dealers may be refraining from fulfilling large customer orders in the when-issued market to avoid running up prices in the auction. Indeed, Bikhchandani et al. (2000) raise the important concern that the when-issued market may be far from being perfectly competitive as well, which calls into question Malvey and Archibald's empirical strategy of utilizing the when-issued market as a competitive benchmark. Moreover, strategic behavior in the when-issued market may be connected with strategic behavior in the auction; thus the assumption that behavior in the when-issued market is invariant to the "downstream" auction mechanism may not necessarily hold.

streams (such as complementary services or access to Fed facilities).

Recovering the marginal valuations, and thus the surpluses of the bidders allows one to construct an upper-bound to the amount of extra revenue that can be derived from switching the auction mechanism – as no voluntary participation mechanism can extract more than the entire surplus that would be obtained in an efficient, surplus-maximizing allocation. We find that the realized total bidder surplus in these auctions amounts to about 3 basis points for the average auction (though surpluses are appreciably higher for Notes auctions than they are for Bills auctions). We also estimate the extent of inefficiency of the allocation to be approximately 2 basis points. These findings suggest that the most cost-savings one can hope for from a redesign of the auction mechanism will be around 5 basis points. Of course, any incentive compatible and individually rational mechanism needs to allow some surplus for the bidders; thus, this is an extremely conservative upper bound.

## 2 Description of our Data Sample and Institutional Background

The data sample used in this study comprises of 975 auctions of Treasury securities conducted between July 2009 and October 2013 (Table 1). The securities in our sample range from 4 week bills to 10 year notes, with 822 auctions of 4-week, 13-week, 26-week, 52-week bills and cash management bills, and 153 auctions of 2-year, 5-year, and 10-year notes. The total volume of issuance through these auctions was 27.3 trillion US dollars, with the average issue size around 28 billion dollars.

The issuance mechanism is a sealed-bid uniform-price auction, which has been the preferred auction mechanism of the Treasury since October 1998. Bids consist of price-quantity schedules and define step functions, with minimum price increments of 0.5 basis points for thirteen, twenty-six, fifty-two week, and cash management bills and 0.1 basis points for all other securities. Noncompetitive, price-taking bids are also accepted but are limited to \$5 million and are usually due before noon, an hour earlier than competitive bids. Noncompetitive bid totals are announced prior to the deadline for competitive tenders.

Bidders in our data are categorized into three major groups: *primary dealers* (PDs), *direct*



*bidders* (DBs), and *indirect bidders* (IBs). During the sample period, 17 to 21 PDs regularly bid in the auctions and made markets in Treasury securities.<sup>4</sup> While we cannot disclose the exact numbers, we can report that there is a lot more IBs than DBs. A Direct Bidder is a non-primary dealer (non-PD) submitter bidding for its own house account, including but not limited to non-PD brokers and dealers, various types of investment funds (pension, hedge, mutual), insurance companies, depository institutions (banks), foreign and international entities (governmental and private), the Federal Reserve (System Open Market Account), and individuals. Typically, most DBs are non-PD brokers and dealers.<sup>5</sup> As for IBs, the largest wealth manager Blackrock is one example<sup>6</sup>, and then (probably) many large pension and mutual funds. PDs can bid on their own behalf (“house bids”) and also submit bids on behalf of the indirect bidders. PDs need to record IBs’ bids on appropriate forms - and hence cannot modify them before submission.

Primary dealers are, as a class of bidders, the largest purchasers of primary issuances. In terms of tendered quantities, primary dealer tenders comprise 69% to 88% of overall tendered quantities. Direct bidders tender 6% to 13% and Indirect bidders 6% to 18% of the tenders.<sup>7</sup> In terms of winning bids, or allocated quantities, we find that Primary Dealers tend to win a smaller proportion of their tendered quantities. Primary Dealers are allocated between 46% to 76% of competitive demand, while Indirect Bidders win the disproportionate share of 17% to 38%, with Direct Bidders’ allocation shares staying close to their tendered quantity shares. Let us now analyze these bidding differences more closely.

---

<sup>4</sup>The list of primary dealers is available here: <https://www.newyorkfed.org/markets/primarydealers>

<sup>5</sup> <http://www.reuters.com/article/us-usa-treasuries-china-idUSBRE84K11720120522> mentions Bank of China and Bank of Japan as direct bidders.

<sup>6</sup><http://www.fa-mag.com/news/club-of-u-s--bond-traders-loses-cachet-in-most-important-market-13848.html>, Bloomberg News, 4/4/2013

<sup>7</sup>Earlier analyses of similar tender and allocation shares by bidder classes have been performed by Garbade and Ingber (2005), Fleming (2007), Fleming and Rosenberg (2007). There have been large changes in the types of bidders participating in and winning securities auctions since 2008, especially for longer maturity securities. In auctions for notes and bonds, Treasury data shows there have been large increases in the proportion of both bids tendered and bids accepted by non-primary dealers. While the typical proportion of indirect bids tendered has remained relatively constant over the time period, Direct Bids tendered by non-primary dealers have increased from virtually none in 2008 to the observed proportions in our data. This has corresponded to a decrease in the proportion of notes and bonds tendered by Primary Dealers.

## 2.1 Analysis of Bid Yields and Bid Quantities

The fact that Primary Dealers are winning a smaller share of their tenders than Direct and Indirect Bidders suggests that Primary Dealers bid systematically higher (lower) yields (prices) in these auctions. To investigate this further, Table 2 reports the quantity-weighted bid-yields submitted by the three bidder groups across different maturities. We include the within-auction standard deviation of quantity-weighted bid yields as a measure of bid dispersion within bidder group. Since bids in these auctions are effectively in the form of demand curves, we also include the total tender quantity submitted by a bidder as a percentage of the issue size (%Issue Size), and the percentage of her tender quantity that the bidder won (%Tender Won).

Looking at the (quantity-weighted) bid yields we see the clear pattern, across all maturities, that Primary Dealers systematically place higher (lower) bid yields (prices) than Direct and Indirect Bidders. The gap in bid yields between Primary Dealers and Indirect Bidders is quite substantial, and ranges between 3 to 18 basis points depending on maturity. Primary Dealers also appear to be bidding 2 to 8 basis points higher yields than Direct Bidders for maturities, except for 2 year notes.

The within-auction dispersion of (quantity-weighted) bid-yields across Primary Dealers is very similar to the dispersion of Direct Bidder bids, ranging from 2 to 7 basis points. Indirect Bidders submit more dispersed bids, especially for the longer-term securities, with the dispersion rising to 19 basis points for 10 year bond auctions.

Primary dealers bid for much larger quantities. The average Primary Dealer offers to purchase between 10% to 20% of issuance, while Direct bidder quantity tenders hover between 3-5% and Indirect Bidders' tenders between 1-3% of the issuance. Given that Primary Dealers tend to bid higher yields, however, it is not surprising that Primary Dealers get allocated a smaller share of their total tenders than Direct or Indirect Bidders. Indeed, while Primary Dealers end up winning only about 20% of their tendered quantities, Direct Bidders win 40-50% and Indirect Bidders, 70-80%.

In Table 3, we accounted for how the quantity-weighted bids compare to the US Treasury published yield (of the respective maturity) on the day of the auction. It is evident that indirect

bidders systematically bid lower yields than the market-level prevailing yield, the direct bidders bid about at the prevailing yields and the primary dealers bid on average above these yields. We next discuss possible drivers of these bid differentials.

## 2.2 Drivers of Bid Differentials?

What might drive these bid differentials across bidder groups? These different bidder groups may have different demand/willingness-to-pay for these securities, depending on their idiosyncratic needs. Some of these bidders are buy-and-hold investors who are replenishing their bond portfolios, while others are broker-dealers whose primary purpose is resale. In particular, it is possible that the reason why Primary Dealers bid lower prices (or equivalently, higher yields) is that they systematically have lower demand/willingness-to-pay for the securities than other bidder classes.

Another possibility is the exercise of market power. Even if bidders' willingness-to-pay is the same for the securities, the Primary Dealers, as we see, are much larger players in this market, commanding significant market share. Such buyers may be able to exercise their monopsony power to try to lower the marginal cost of acquisition.

A particularly detrimental source of market power leading to higher yields bids, of course, is collusion among primary dealers. Following the recent settlements in the LIBOR fixing scandal and the Foreign Exchange rate fixing case, there have, indeed, been a number of recent allegations regarding collusive behavior among primary dealers in the treasury auctions as well.<sup>8</sup> However, the raised allegations regarding misconduct have been about information sharing between the primary dealers; one particular allegation is that primary dealers communicated about the amount of customer/indirect bidder interest they were receiving prior to the auction. Such information sharing does not necessarily lead to coordinated conduct; indeed, sharing information about private values/costs prior to an auction can lead to potentially stiffer competition a la Bertrand compared to the private information case.<sup>9</sup> Moreover, as reported in footnote 3 above, we find that the auction prices are higher than when-issued market prices. If dealers were colluding in the auction after sharing information, they would be driving auction prices down. Nevertheless, our model, which

---

<sup>8</sup>See, for example, <http://www.bloombergvew.com/articles/2015-09-18/maybe-banks-manipulated-treasury-bonds-too>.

<sup>9</sup>Moreover, in a common values context, the sharing of information may lead to a reduction in the winner's curse, which may also lead to an increase in auction revenue, as pointed out by Boyarchenko, Lucca and Veldkamp (2016).

we describe below, can be tailored to take into account the precise nature of such information sharing arrangements.

Yet another possibility for the high-yield/low-price bids by the Primary dealers is the fact that they face minimum bid constraints. The Treasury expects primary dealers to bid in every auction, and to bid, at a minimum, for their pro-rata share of the auction volume based on the number of primary dealers at the time of the auction. However, in our data, we find primary dealers to be bidding quantities close to their pro-rata share only very infrequently – in our entire data set, there are only 23 instances (out of 19,000 bids) in which a primary dealer bid for a quantity that is within 1% of its minimum bidding requirement. This suggests that the minimum bidding requirement is only very rarely binding.

Table 4 investigates the differentials in bid yields implied by Table 2 through regressions. We have split the sample into auctions of Treasury Bills and Treasury Notes, as it is possible that the market dynamics are very different across these different classes of securities. We also control for auction fixed effects in each regression; thus the regressions provide within-auction comparisons that account for differing supply-demand conditions that affect the level of the bids.

The first and third specifications regress (quantity-weighted) bid yield on indicators for Direct and Indirect Bidders in the Bills and Notes sectors. We find here the pattern implied by Table 2: Primary Dealers systematically (and statistically significantly) bid higher yields than Direct and Indirect Bidders. Primary Dealer bids are 2 (4) basis points higher than Direct (Indirect) bids in the Bills sector, and 6 (11) basis points higher in the Notes sector.

The second and fourth columns of Table 4 include the bidder's share of the total tender size as a proxy for bidder size. There are two main ways through which bidder size may affect the bids: bidders demanding larger quantity may have higher demand for the security, but they may also have higher market power. The regressions indicate that larger bidders systematically bid higher yields. The effect is quite large – the coefficient estimate indicates that a size increase of 10% of total issue size accounts for 1(6) basis point increase in the bid yield.

Accounting for bidder size appears to lower the differences in bid yields across bidder classes, but Primary Dealers appear to bid higher yields than Direct and Indirects even accounting for their

demanded share of the total issue.

As we noted above, since bids reflect both differences in demand and also differences in market power, it is difficult to interpret these documented differences in bids. Even though we find that larger bidders bid higher yields, this is not *prima facie* evidence that large bidders exercise market power; it is possible that larger bidders also have lower demand.

In the next section, we will describe a model of bidding that will allow us to separate out the market power and demand components of bid heterogeneity. The model, and the measurement it will allow us to conduct, will *rely on the assumption of bidder optimization*. In essence, what we will end up doing is to measure the elasticity of (expected) residual supply faced by each bidder. This is directly observable in the data, and does not require behavioral assumptions. This elasticity will be our measure of the “potential market power” possessed by each bidder. Assuming that bidders are expected profit maximizers who will exercise their market power in a unilateral, noncooperative fashion, we can then estimate the willingness-to-pay/demand that rationalizes the observed bid.

### 3 Model of Bidding

Our analysis is based on the share auction model of Wilson (1979), and Kastl (2011), which allows for discrete bids. In Hortaçsu and Kastl (2012), we further adapted this model to allow primary dealers to observe the bids of others, hence allowing for “indirect bidders,” whose bids are routed by primary dealers.

There are three classes of bidders:  $N_P$  *primary dealers* (in index set  $\mathcal{P}$ ),  $N_I$  potential *indirect bidders* (in index set  $\mathcal{I}$ ) and  $N_D$  potential *direct bidders* (in index set  $\mathcal{D}$ ). They are each bidding for a perfectly divisible good of (random)  $Q$  units. We assume that the number of *potential* bidders of each type participating in an auction,  $N_P, N_I, N_D$ , is commonly known. However, except primary dealers, the exact number of indirect and direct bidders is not known. Before bidding commences, bidders observe private (possibly multidimensional) signals. Let us denote these signals for the different bidder groups as  $S_1^P, \dots, S_{N_P}^P, S_1^I, \dots, S_{N_I}^I, S_1^D, \dots, S_{N_D}^D$ . The bidding then proceeds in two stages. In stage 1, indirect bidders submit their bids to their primary dealer. These bids, denoted by  $y^I(p|S_j^I)$ , specify for each price  $p$ , how big a share of the securities offered in the auction indirect

bidder  $j$  demands as a function of her private information  $S_j^I$ . In stage 2, direct bidders submit their bids,  $y^D(p|S_j^D)$ , and primary dealer  $k$  submits her customers' bids, and also places her own bid,  $y^P(p|S_k^P, Z_k^P)$ , where  $Z_k^P$  contains all information dealer  $k$  observes from seeing the bids of its customers. If dealer  $k$  observes the bids of not only her customers but also of other dealers', as has been recently alleged, one can appropriately modify the information set of this dealer. Winning  $q$  units of the security is valued according to a marginal valuation function  $v_i(q, S_i)$ , which we assume is symmetric within each class of bidders. We further assume:

**Assumption 1**  $v^g(q, S_i^g)$  is non-negative, bounded, strictly increasing in (each component of)  $S_i^g$   $\forall q$  and weakly decreasing in  $q \forall s_i^g$ , for  $g \in \{P, I, D\}$ . Direct and indirect bidders' and dealers' private signals are independent and drawn from a common support  $[0, 1]^M$  according to atomless distribution functions  $F^P(\cdot)$ ,  $F^I(\cdot)$ , and  $F^D(\cdot)$ .

Under this assumption, we have a setting with *private*, not interdependent values, as in Chapman, McAdams and Paarsch (2007), Kang and Puller (2008), Hortaçsu and McAdams (2010), Kastl (2011).<sup>10</sup> An interdependent/common value element may be present in the Treasury auction context, as resale is a primary motivation for many of the bidders. However, a common value environment arises when bidders have private information about an underlying common value for the security, so that learning another bidder's information modifies the valuation other bidders attach to the security. This is an environment where bidders (especially primary dealers) are sophisticated financial traders who all have access to the same macroeconomic and financial data to assess the value of the security at the time of the auction. Moreover, the primary auction market is preceded by an active when-issued market, which was designed explicitly to aggregate the dispersed private information of the bidders. We thus assume that the heterogeneity in values that remains at the time of the auction can be largely attributed to private values that stem from heterogeneities in investment and inventory needs/opportunities across bidders.

Ideally, one should explicitly test for the null hypothesis of private values vs. the alternative of common values. In Hortaçsu and Kastl (2012), we utilized data on Canadian Treasury auctions

---

<sup>10</sup>Strictly speaking, independence is not necessary for our characterization of equilibrium behavior in this auction, but we impose it in our empirical application.

on updates in primary dealers' bids in response to indirect bids they routed. Under the null hypothesis of private values, observed indirect bids provide information only about the distribution of residual supply, whereas under common/interdependent values, indirect bids would also allow primary dealers to update beliefs about the unknown common value of the security. We were not able to reject the null hypothesis of private values for 3 month and 12 month Canadian Treasury securities. While utilizing a similar testing strategy for the U.S. context would be very desirable, data on "updates" of primary dealer bids is not available for the present study, and thus left for future research.

Let  $\theta_j$  denote private information of bidder  $j$ , i.e., for a direct bidder  $\theta_j \equiv S_j^D$ , indirect bidder  $\theta_j \equiv S_j^I$  and for a primary dealer  $\theta_j \equiv (S_j^P, Z_j^P)$ . Bidders' pure strategies are mappings from private information in each stage to bid functions  $\sigma_i : \Theta_i \rightarrow \mathcal{Y}$ , where the set  $\mathcal{Y}$  includes all admissible bid functions.

The expected utility of type  $\theta_i$ -bidder (from group  $g \in \{P, D, I\}$ ) who employs a strategy  $y_i^g(\cdot|\theta_i)$  in a uniform price auction given that other bidders are using  $\left\{y_j^{(g,-g)}(\cdot|\cdot)\right\}_{j \neq i}$  can be written as:

$$\begin{aligned} EU_i^g(\theta_i) &= \mathbb{E}_{Q, \Theta_{-i}|\theta_i} u^g(\theta_i, \Theta_{-i}) \\ &= \mathbb{E}_{Q, \Theta_{-i}|\theta_i} \left[ \int_0^{Q_i^c(Q, \Theta, \mathbf{y}^{(g,-g)}(\cdot|\Theta))} v_i^g(u, \theta_i) du - P^c(Q, \Theta, \mathbf{y}^{(g,-g)}(\cdot|\Theta)) Q_i^c(Q, \Theta, \mathbf{y}^{(g,-g)}(\cdot|\Theta)) \right] \end{aligned}$$

where  $Q_i^c(Q, \Theta, \mathbf{y}^{(g,-g)}(\cdot|\Theta))$  is the (market clearing) quantity bidder  $i$  obtains if the state (bidders' private information and the supply quantity) is  $(Q, \Theta)$  and bidders bid according to strategies specified in the vector  $\mathbf{y}^{(g,-g)}(\cdot|\Theta)$ . Similarly  $P^c(Q, \Theta, \mathbf{y}^{(g,-g)}(\cdot|\Theta))$  is the market clearing price associated with state  $(Q, \Theta)$ . This is the expected consumer surplus, as given by the expected area under the demand curve up to the random allocation,  $Q_i^c$ , minus the expected payment, which depends on the random allocation and random market clearing price,  $P^c$ .

Our solution concept will be *group symmetric Bayesian Nash Equilibrium*, in which direct and indirect bidders submit bid functions that are symmetric up to their private signals, i.e.  $y_j^D(p|S_j^D) = y^D(p|S_j^D)$ ,  $j \in \mathcal{D}$ , and  $y_j^I(p|S_j^I) = y^I(p|S_j^I)$ ,  $j \in \mathcal{I}$ .<sup>11</sup> Bidders' choice of bidding strategies is restricted to non-increasing step functions with an upper bound on the total

<sup>11</sup>Conditions for existence of Bayesian Nash Equilibria are explored in Kastl (2012).

quantity they can win (up to 35% of the total quantity). When bidders use step functions as their bids, rationing occurs except in very rare cases. We will thus assume, as it is in practice, *pro-rata on-the-margin* rationing, which proportionally adjusts the marginal bids so as to equate supply and demand. Also, in extremely rare situations where multiple prices clear the market (due to discreteness of quantities), we assume that the auctioneer selects the highest market clearing price.

### 3.1 Characterization of equilibrium bids

The key source of uncertainty faced by the bidders in the auction is the market clearing price,  $P^c$ , which maps the state of the world,  $(\mathbf{s}^I, \mathbf{s}^D, \mathbf{s}^P, \mathbf{z})$  into prices through equilibrium bidding strategies.

Let us now define the probability distribution of the market clearing price from the perspective of a direct bidder  $j$ , who is preparing to make a bid  $y^D(p|s_j)$ . The probability distribution of the market clearing price from the perspective of direct bidder  $j$  will be:

$$\Pr(p \geq P^c | s_j) = \mathbb{E}_{\{S_{k \in \mathcal{D} \cup \mathcal{P} \cup \mathcal{I} \setminus j}, Z_{l \in \mathcal{P}}\}} \mathbb{I} \left( Q - \sum_{m \in \mathcal{P}} y^P(p|S_m, Z_m) - \sum_{l \in \mathcal{I}} y^I(p|S_l) - \sum_{k \in \mathcal{D} \setminus j} y^D(p|S_k) \geq y^D(p|s_j) \right) \quad (1)$$

where  $E_{\{\cdot\}}$  is an expectation over all other bidders' (including indirect bidders, primary dealers, and other direct bidders) private information, and  $I(\cdot)$  is the indicator function. In words, the probability that the market clearing price  $P^c$  will be below a given price level  $p$  is the same as the probability that the (uncertain) residual supply of the security at price  $p$  will be higher than the quantity demanded by bidder  $j$  at that price.

Primary dealers can condition their bids on the indirect bidders' bids they observe, thus this expression will be different. 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. The distribution of the market clearing price from the perspective of primary dealer  $j$ , who observes the bids submitted by indirect bidders  $m$  in an index set  $\mathcal{M}_j$ , is  $\Pr(p \geq P^c | s_j, z_j)$ , and equal to:

$$\mathbb{E}_{\{S_{k \in \mathcal{I} \setminus \mathcal{M}_j}, S_{l \in \mathcal{D}}, S_{n \in \mathcal{P} \setminus j}, Z_{n \in \mathcal{P} \setminus j} | z_j\}} \mathbb{I} \left( Q - \sum_{k \in \mathcal{I} \setminus \mathcal{M}_j} y^I(p|S_k) - \sum_{l \in \mathcal{D}} y^D(p|S_l) - \sum_{n \in \mathcal{P} \setminus j} y^P(p|S_n, Z_n) \geq y^D(p|s_j, z_j) + \sum_{m \in \mathcal{M}_j} y^I(p|s_m) \right)$$

The main difference in this equation compared to equation (1) is that the dealer conditions on all



observed customers' bids, all bids in index set  $\mathcal{M}_j$ . This is exactly where 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.<sup>12</sup>

An indirect bidder recognizes that her bid will be observed by a primary dealer,  $m$ , and can condition on the information that she provides to this dealer. Thus, for indirect bidder  $j$ , who submits her bid through a primary dealer  $m$ , we have:

$$\Pr(p \geq P^c | s_j) = \mathbb{E}_{\{S_k \in \mathcal{I} \setminus j, S_l \in \mathcal{D}, S_n \in \mathcal{P}, Z_n \in \mathcal{P} | s_j\}} \mathbb{I} \left( Q - \sum_{k \in \mathcal{I} \setminus j} y^I(p | S_k) - \sum_{l \in \mathcal{D}} y^D(p | S_l) - \sum_{n \in \mathcal{P}} y^P(p | S_n, Z_n) \geq y^I(p | s_j) \right)$$

where  $y^I(p | s_j) \in Z_m$ .

The necessary condition for optimal bidding is given by:

**Proposition 1** (*Kastl 2012*) *Under assumption 1, in any Bayesian Nash Equilibrium of a Uniform Price Auction, for a bidder of type  $\theta_i$  submitting  $\hat{K}(\theta_i)$  steps, every step  $(q_k, b_k)$  characterizing the equilibrium bid function  $y(\cdot | \theta_i)$  has to satisfy:*

$$v_i(q_k, \theta_i) = \mathbb{E}(P^c | b_k > P^c > b_{k+1}, \theta_i) + \frac{q_k}{\Pr(b_k > P^c > b_{k+1}, \theta_i)} \frac{\partial \mathbb{E}(P^c; b_k \geq P^c \geq b_{k+1}, \theta_i)}{\partial q_k} \quad (2)$$

$\forall k \leq \hat{K}(\theta_i)$  such that  $v(q, \theta_i)$  is continuous in a neighborhood of  $q_k$ .

This expression is very close to  $MC = E[P(Q)] + E[P'(Q)]Q$ , i.e., to an oligopolist’s optimality condition in a setting where the oligopolist faces uncertain demand in the spirit of Klemperer and Meyer (1989). Kastl (2011) points out that bids *above* marginal values may be optimal in a uniform price auction with restricted strategy sets. The intuition has to do with the restriction on the strategies requiring the bidders to “bundle” bids for several units together and thus to trade-off potential (ex post) loss on the last unit in the bundle against the probability of obtaining the high-valued infra-marginal units in the bundle.<sup>13</sup>

<sup>12</sup>If primary dealers engage in information sharing about their customer bids, as has been alleged recently, the exact nature of this sharing arrangement can be reflected in this optimality condition.

<sup>13</sup>For example, consider one very small bidder, so that he is a “price taker.” Assume also a non-degenerate distribution of the market clearing price with continuous and strictly positive density over a compact support and let  $K_i = 1$ ; i.e. the bidder is constrained to submit a single step as her bid. In this case, the second term on the

For our empirical exercise, we define the notion of bid-shading based on the above condition for optimal bidding. Typically, bid-shading is defined as the difference between a bidder’s value and her bid. As highlighted above, however, this notion of shading might often result in negative values in auctions where bids are constrained to be step functions.

**Definition 1** *The average bid shading is defined as:* 
$$B(\theta_i) = \frac{\sum_{k=1}^{K_i} q_k [v_i(q_k, \theta_i) - b_k]}{\sum_{k=1}^{K_i} q_k}.$$

This is the quantity-weighted average shading on the marginal (i.e., last) infinitesimal unit demanded at step  $k$ . A negative value does not imply negative surplus, as surplus can be positive on the inframarginal units. A negative value is due to the combination of two factors: (i) the perceived market power of this bidder at  $k^{th}$  step is probably small and (ii) this bidder believes that if  $k^{th}$  step were marginal, the market clearing price will likely be much lower than her bid.

Given equation (2), a definition of shading that only takes positive values is the following:

**Definition 2** *The average shading factor is defined as:* 
$$S(\theta_i) = \frac{\sum_{k=1}^{K_i} q_k [v_i(q_k, \theta_i) - \mathbb{E}(P^c | b_k > P^c > b_{k+1}, \theta_i)]}{\sum_{k=1}^{K_i} q_k}.$$

This is a quantity-weighted measure of shading, where shading at step  $k$  is defined as the difference between the marginal value,  $v_i(q_k, \theta_i)$  and the expected market clearing price, conditional on  $k^{th}$  step being marginal,  $\mathbb{E}(P^c | b_k > P^c > b_{k+1}, \theta_i)$ . This measure of shading is non-negative, since market clearing price is non-decreasing in quantity demanded. Another way to interpret this shading factor is to note that it corresponds to the weighted sum of the second term on the right-hand side of equation (2), which is essentially the expected inverse elasticity of the residual supply curve faced by the bidder.

## 4 Estimating Marginal Valuations

To estimate the rationalizing marginal valuations, we use the “resampling” method developed in Hortaçsu (2002), Kastl (2011), and Hortaçsu and Kastl (2012). The details and the asymptotic behavior of our estimator is described in detail in Hortaçsu and Kastl (2012) and Cassola, Hortaçsu

---

RHS of (2) vanishes because of the bidder being a price taker. This bidder thus optimally asks for a quantity such that his marginal valuation at that quantity is equal to the expected price conditional on this price being lower than his bid, i.e.  $v_i(q_k, \theta_i) = \mathbb{E}(P^c | b_k > P^c, \theta_i)$ . Therefore, whenever there is a positive probability of the market clearing price being below his bid, his bid will be higher than his marginal valuation for that quantity.

and Kastl (2013). The “resampling” method that we employ is to draw from the empirical distribution of bids to simulate different realizations of the residual supply function that can be faced by a bidder, thus obtaining an estimator of the distribution of the market clearing prices. In equation (2), the only other object we need to estimate is the slope of this expectation, which we estimate using a standard numerical derivative as in Kastl (2011).

In the present case, we have three classes of bidders:  $N_P$  *primary dealers* (in index set  $\mathcal{P}$ ),  $N_I$  potential *indirect bidders* (in index set  $\mathcal{I}$ ) and  $N_D$  potential *direct bidders* (in index set  $\mathcal{D}$ ). To estimate the probability in equation (1) for direct and indirect bidders, we draw direct and indirect bids from the empirical distribution of these classes of bids (we augment the data with zero bids for non-participating direct and indirect bidders). Clearly, PDs bids are not necessarily independent of IBs’ bids. In particular, a given PD’s bid is a function of IBs’ bids that he/she observes. Hence, we begin each iteration by drawing an indirect bidder and employ kernel-weighting to draw a PD’s bid: for a given vector of drawn IBs’ bids, draw only from those PDs’bids that were submitted with similar information as measured by the kernel.

Also, to estimate the probability distribution from the perspective of primary dealers, we need to take into account the full information set of the dealer. This is achieved by a slight modification of the above procedure: fixing a primary dealer, who has seen  $M$  indirect bids, we draw  $N_I - M$ , rather than  $N_I$ , indirect bids, and take the observed indirect bids along with the dealer’s own bid as given when calculating the market clearing price. I.e., we subtract the *actual observed* customer bid from the supply before starting the resampling procedure.

Unobserved heterogeneity across auctions that may be driving valuations is a big concern in the empirical auctions literature. The danger this creates is the potential pooling of bidding data across auctions that have very different demand structures, which may cloud inference regarding the probability distribution of the market clearing price. Another, related, concern is the potential for multiple strategic equilibria – bidders may be playing different equilibria in different auctions in the data set. To combat these issues, we use marginal valuations auction-by-auction; using data on bids from only one auction at a time. While this reduces precision of our estimates, the volatile economic environment especially in 2009 and 2010 suggests that auctions of the same security

at different times may be subject to very different demand-side factors and that accounting for unobservables at the auction level may be very important. We discuss the consistency property of the single-auction estimation scheme in Cassola et al. (2013).<sup>14,15</sup>

## 5 Results

### 5.1 Bid Shading Analysis

As we discussed in our analysis of bids, the difference in bids across bidder groups may arise from two separate factors: differential ability to exercise market power, i.e. bid shading, vs. differential willingness-to-pay for the issued security. Our estimation method yields estimates of the two terms on the right hand side of Equation (2) based on the empirical distribution of bids within each auction in our data set. Using these, we can construct an estimate of the marginal valuation for each bid step, which can then be utilized to compute the two different shading factors we defined in the previous section for each bidder and auction.

Table 5 reports the results of regressions similar to those for bids. The first two columns of the table look at the differences in bid shading (according to Definition 1 above) across bidder groups for the Bill sector. Column (1) implies that Primary Dealers shade their bids 1.9 basis points more than Direct Bidders, and 3.5 basis points more than Indirect Bidders. Column (2) introduces the bidder size control, and we find that the shading differentials decline slightly, to 1.4 basis point against Direct Bidders and 3 basis points against Indirect Bidders. We also find, intuitively, that larger bidders choose to shade their bids more. The coefficient estimate suggests that going from zero to 10% market share is associated with 0.3 basis points more bid shading.

Columns (3) and (4) repeat the same analysis for Definition 2 of bid shading, and find qualita-

---

<sup>14</sup>What is needed for a non-degenerate distribution of the market clearing price as the number of participants within an auction grows large (and thus also for non-degenerate strategy in the limit game) is uncertainty in the available supply that does not vanish. This can come in our application, for example, from the uncertain demand by foreign direct bidders (Bank of China or Bank of Japan).

<sup>15</sup>Utilizing auction-by-auction estimation also allows our estimates to account for the recent allegation that primary dealers may have been manipulating the forward/when issued-price (see, e.g. <http://www.bloombergvew.com/articles/2015-09-18/maybe-banks-manipulated-treasury-bonds-too>). Variation across time in the severity of such a manipulation scheme would affect an auction level variable observed by bidders, but not necessary the econometrician, who may take the observed when-issued price at face value. Auction-by-auction estimation, however, allows the econometrician to account for this source of unobserved heterogeneity across auctions.

tively the same result.

Columns (5) through (8) repeat the analysis for the Notes sector. Here, we see even larger differentials in shading. Primary Dealers shade their bids (according to Definition 1) 5 basis points more than Direct Bidders and 13 basis points more than Indirect Bidders. Putting in the control for bidder size, once again we find that larger bidders can shade their bids more: going from zero to 10% market share increases shading by 3 basis points. The size control diminishes the shading differential between Primary Dealers and Direct and Indirect Bidders (to 2 and 10 basis points), but, once again, does not eliminate the differential. Once again, Columns (7) and (8) repeat the analysis with the second definition of bid shading. The qualitative results are the same as for the Definition 1.

The reader may be worried that bidders may submit “artificial” bids so as to satisfy bidding requirements, but such bids may be at prices that have no chance of clearing the market (for example, bidding a yield of 90%), and our measure of quantity-weighted shading might be sensitive to such a behavior. We, therefore, repeat the analysis of bid shading from Table 5, but we now focus on shading at the marginal step. The results are reported in Table 6. All qualitative features of the results (significant, but modest shading and its relative ordering between the various bidder groups) remain unchanged. Recall from the discussion on page 17 that the bid shading at the marginal quantity may be negative due to discreteness of bidding.

### **Demand Differentials or Bid-Shading Differentials?**

Recall that our analysis of bids in Table 4 revealed that, controlling for size, Primary Dealers bid 1 (2.5) basis points higher yields than Direct(Indirect) bidders for Bills, and 1 (4) basis points higher yields for Notes. Since we found that Primary Dealers shade their bids 1.4 (3) basis points more than Direct (Indirect) for Bills, and 2 (10) basis points more than Direct (Indirect) bidders for Notes, the bid differentials are rationalized by Primary Dealers having 0.4 (0.5) basis points higher willingness-to-pay for Bills, and 1 (6) basis points higher willingness-to-pay for Notes – again, controlling for bidder size. Therefore, our results suggest that, under the assumption of expected profit maximization, the main reason why Primary Dealers bid higher yields than other bidder groups is not because they have lower valuation for the securities, but because they are able to

exercise more market power.

We can gain some intuition on this finding from considering a thought experiment of a bidder submitting a bid against the expected residual supply curve. If this bidder expects that she might be marginal, i.e. that her bid can affect the market clearing price, then she should optimally shade her bid. This is the same trade-off that an oligopolist faces when solving a profit maximization problem. If the residual supply were flat, i.e. bidder knew with certainty that he cannot affect the market clearing price, he would simply bid his value. Whenever the expected residual supply has some slope, she would, however, shade her bid. Therefore, the key source of variation that identifies the source of heterogeneity is the (average) shape of the residual supplies that a bidder may face around the point where they intersect this bidder's bid. This is driven by shapes of other bidders' bids. For example, if all bids were identical (and the auction were oversubscribed), clearly no bidder would be able to affect the market clearing price individually and hence everyone would be bidding their values. The lower the elasticity of the expected residual supply, the larger the markup: and we find that the primary dealers face more inelastic residual supply curves.

## 5.2 Infra-marginal Surplus Analysis

We now quantify how much infra-marginal surplus bidders are getting from participating in these auctions. We can utilize Equation (2) to calculate the bidders' marginal valuations, and use these to compute *ex-post surplus* each bidder gains on the units that they win in the auction. We obtain point estimates of the “rationalizing” marginal valuation function  $v(q, s)$  at the (observed) quantities that the bidders request. We then compute the area under the upper envelope of the inframarginal portion of the marginal valuation function, and subtract the payment made by each bidder.

We should provide abundant caution regarding what “infra-marginal bidder surplus” means. Any (counterfactual) auction system has to allow bidders to retain some surplus. Indeed, Figure 1 shows that even if bidders bid perfectly competitively, i.e. reveal their true marginal valuations without any bid shading, they would gain some surplus from the auction, just because they have downward sloping demand curves. Indeed, if there are any costs of participating in the auction, it

would have to be justified by the expected surplus. In terms of assessing the cost effectiveness of the issuance mechanism, the most we can say is that the bidder surplus under an efficient allocation reflects a conservative upper bound to the amount of cost-saving that can be induced by a change in issuance mechanism.

Table 7 reports the infra-marginal surpluses enjoyed by different bidders groups across the maturity spectrum. We report the surpluses in basis points, and also report the total infra-marginal surpluses accrued to the bidders during our sample period of July 2009 to October 2013. Direct and Indirect bidder surpluses are between 0.02 and 3.58 basis points across the maturity spectrum, with the shorter end of the maturity spectrum generating very low surpluses in general. Once again, these surpluses may reflect the outside option of not buying these securities in auction and purchasing them in the when-issued or resale markets – and appear sensible given the differentials between auction prices and secondary market rates. Aggregating the surpluses over the entire set of auctions in our data set (which amounted to about \$27 trillion in issue size), we find Direct and Indirect Bidders’ aggregate surplus to be about \$1.6 billion, or about 0.6 basis points.

Primary Dealers’ infra-marginal surplus, however, appears to be significantly larger. For Primary Dealers, the derived surplus might not necessarily be in line with the differentials with the quoted secondary market prices of these securities. Primary Dealers’ demand is typically quite large, and fulfilling such levels of demand is likely to have a price impact in the secondary markets. Moreover, retaining Primary Dealership status has a number of complementary value streams attached to it beyond the profits derived from reselling the new issues. For example, being a Primary Dealer allows firms access to open market operations and, especially in this period, the QE auction mechanism that is exclusive to primary dealerships. Between March 2008 and February 2010, Primary Dealers also had access to a special credit facility from the Fed to help alleviate liquidity constraints during the crisis. In contrast, one may expect the surpluses attained by Direct and Indirect bidders to be driven more by their outside options of purchasing these securities from secondary markets.

We find that Primary Dealers derive most of their infra-marginal surplus from the longer end (2 to 10 year notes) of the maturity spectrum. There may be a number of reasons why demand for this

part of the maturity spectrum is more heterogeneous across bidders. One possibility is the presence of different portfolio needs across dealers' clientele. Moreover, there are typically alternative uses for such securities beyond simple buy-and-hold – Duffie (1996) shows that this part of the spectrum can be particularly valuable for its use as collateral in repo transactions. Surpluses derived from the shorter end of the maturity spectrum, which may have fewer alternative uses, are much smaller. Overall, we find that Primary Dealers' derived surplus aggregated to \$6.3 billion during our sample period. Compared against the \$27 trillion in issuance, Primary Dealer surplus makes up for 2.3 basis points of the issuance. Along with the Direct Bidder and Indirect Bidder surpluses, we find that bidder surplus added up to 3 basis points during this period.

Once again, we should emphasize that any other issuance mechanism would have to provide bidders with surpluses to ensure participation and to reward them for their private information. Moreover, even if bidders are behaving in a perfectly competitive manner, without displaying any bid-shading, they would enjoy surpluses. However, we can conservatively estimate that revenue gains from further optimizing the issuance mechanism is bounded above by 3 basis points.

Table 8 runs regressions on the calculated bidder surpluses that closely resemble those in Tables 4 and 5. The surpluses here are reported in thousands of dollars, and all regressions control for auction fixed effects, giving us within auction comparisons. In Column (1), we find that Direct and Indirect bidders gain significantly lower surpluses than Primary Dealers (the excluded category appearing in the constant), and that Indirect bidder surplus especially is not statistically different from zero. Column (2) adds in the bidder size control, measured as the bidder's tender size as percentage of total supply (% Q Total). We find that larger bidders indeed gain higher surpluses. However, Direct and Indirect Bidders gain lower surpluses than Primary Dealers even when size is controlled for.

Column (3) introduces a new control variable – we have added here the number of Indirect Bidders whose bids a Primary Dealer routes in the auction. This variable is a rough proxy for the order-flow information that the Primary Dealer is privy to. Indeed, the regression reveals a significant correlation between the number of Indirect Bidders who routed their bids through a Primary Dealer, and the surplus (controlling for the bidder's size). An additional Indirect Bidder



going through a Primary Dealer would be correlated with about seven thousand dollars more in surplus.

Columns (4) through (6) repeat the same analysis for the Notes sector. Measured Primary Dealer surplus in this sector (which is measured by the constant term in our regression) is much larger as compared to their surplus in Bills auctions. Direct and Indirect Bidders gain much smaller surpluses compared to the Primary Dealers – indeed, Indirect Bidder surplus is very close to zero.

When we control for bidder size in Column (5), we find a very large benefit to being large. An increase in market share from 0 to 10% of the issue size is correlated with a rise in surplus of \$830k in the Notes auctions. Once again, though, we should stress that this is not necessarily due to market power. It is very possible that larger bidders also have higher demand, and thus derive more surplus from the auctions.

Finally, we introduce the number of Indirect Bidders routed by Primary Dealers in Column (6). Here, we find that each additional Indirect Bidder observed is associated with a \$200K gain in Primary Dealer surplus. Since Primary Dealers on average route 2.5 Indirect Bids in Notes auctions, this estimate suggests that we can ascribe about \$500K or about 25% of their surplus in Notes auctions to information contained in Indirect bids. However, there are important caveats to interpreting this as the “value of order flow.” It is possible that Primary Dealers who observe more Indirect bids may have systematically higher valuations for the securities, and hence may be getting higher surpluses due to this.<sup>16</sup>

### 5.3 Efficiency

Using our estimates of marginal values, we can compute the efficient surplus. Fixing the supply in an auction, we can construct a measure of efficiency by comparing the surplus from the efficient allocation to the surplus that is achieved in the actually implemented allocation. Our efficiency estimates are reported in Table 9. Overall, the efficiency losses seem to be quite modest, amounting to just over 2 basis points on average. Especially on the shorter side of the maturity spectrum the

---

<sup>16</sup>In our prior work on Canadian Treasury auctions (Hortaçsu and Kastl (2012)) we focused on revisions of primary dealers upon getting indirect bid information. This way, we were able to conduct a “but-for” analysis of how indirect bids contribute to primary dealers’ surplus. Unfortunately, we do not observe bid revisions – though in the Canadian context we found that customer/indirect bids made up between 13-27% of primary dealer surplus, which is close to the 25% estimate we are getting here.

auctions are quite efficient. In auctions of bills, only rarely does the loss exceed 1 basis point. This is likely a consequence of bidders submitting very similar (and flat) bids in these short-term auctions. Therefore, the extent of misallocation is limited. On the other hand, at long maturities, the efficiency loss is slightly higher - up to 6.4 basis points in 10-year notes auctions. The larger efficiency loss here is driven by the much larger heterogeneity of expressed bids in these auctions.<sup>17</sup>

## 6 Conclusion

Using our data on U.S. Treasury auctions conducted between July 2009 and October 2013, we have documented significant differences in bidding behavior across the three different bidder groups: Primary Dealers, Direct Bidders, and Indirect Bidders. We provide a model to decompose the bidding differentials into differences in demand/willingness to pay, and differences in ability to exercise market power. We estimate market power by assuming optimizing behavior and rational expectations about the elasticity of residual supply. Our results suggest that opportunities to exercise market power do exist in this market, and that Primary Dealers especially have the potential to shade their bids significantly – to the extent that their bids are lower (higher yield) than others, even though their willingness-to-pay is higher.

We also quantify the bidder surpluses that rationalize observed bids within our model. We estimate total bidder surplus to be about 3 basis points of the total issue size, with higher surpluses in Treasury Notes auctions as compared to Treasury Bills auctions. Since we estimate that the efficiency loss from the allocation of bills or notes to bidders with relatively lower values is around 2 basis points, 5 basis points is a conservative upper bound on the amount of cost-savings that the Treasury can gain from optimizing its issuance mechanism.

---

<sup>17</sup>This may stem from the fact that there is more uncertainty in long-term macroeconomic forecasts, and that bidders may have more heterogeneous uses for these securities (e.g. buy-and-hold, resale, repo collateral) than for short-term securities.

## References

- 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.
- Bikhchandani, Sushil, Patrik Edsparr, and Chi fu Huang**, “The Treasury bill auction and the when-issued market: Some evidence,” *Decisions, Operations, and Technology Management*, 2000.
- Boyarchenko, Nina, David Lucca, and Laura Veldkamp**, “Taking Orders and Taking Notes: Dealer Information Sharing in Financial Markets,” August 2016. NBER wp 22461.
- 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.
- Chapman, James T. E., David McAdams, and Harry J. Paarsch**, “Bounding Revenue Comparisons across Multi-Unit Auction Formats under  $\varepsilon$ -Best Response,” *The American Economic Review*, 2007, *97* (2), pp. 455–458.
- Duffie, Darrell**, “Special Repo Rates,” *The Journal of Finance*, 1996, *51* (2), pp. 493–526.
- Elyakime, Bernard, Jean-Jacques Laffont, Patrice Loisel, and Quang Vuong**, “Auctioning and Bargaining: An Econometric Study of Timber Auctions with Secret Reservation Prices,” *Journal of Business & Economic Statistics*, April 1997, *15* (2), pp.209–220.
- Fleming, Michael**, “Who Buys Treasury Securities at Auction?,” *Current Issues in Economics and Finance*, 2007, *13*, Federal Reserve Bank of New York.
- and **Joshua V. Rosenberg**, “How Do Treasury Dealers Manage Their Positions?,” August 2007. Federal Reserve Bank of New York Staff Reports, No. 299.
- Friedman, Milton**, *A Program For Monetary Stability*, Fordham University Press, 1960.

- Garbade, K. D.**, “The Institutionalization of Treasury Note and Bond Auctions, 1970-75,” *Federal Reserve Bank of New York Economic Policy Review*, May 2004.
- Garbade, Kenneth D.**, “Why the U.S. Treasury Began Auctioning Treasury Bills in 1920,” *Federal Reserve Bank of New York Economic Policy Review*, 2008, *14*, Federal Reserve Bank of New York.
- and **Jeffrey Ingber**, “The Treasury Auction Process: Objectives, Structure, and Recent Adaptations,” *Federal Reserve Bank of New York Current Issues in Economics and Finance*, 2005, *11*, Federal Reserve Bank of New York.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong**, “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 2000, *68* (3), pp. 525–574.
- 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.
- Jegadeesh, Narasimhan**, “Treasury Auction Bids and the Salomon Squeeze,” *The Journal of Finance*, Sep 1993, *48* (4), pp.1403–1419.
- Kang, Boo-Sung and Steven L. Puller**, “The Effect of Auction Format on Efficiency and Revenue in Divisible Goods Auctions: A Test Using Korean Treasury Auctions,” *Journal of Industrial Economics*, 2008, *56* (2), pp. 290–332.
- 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.

**Klemperer, Paul and Margaret Meyer**, “Supply function equilibria in oligopoly under uncertainty,” *Econometrica*, 1989, 57 (6), pp.1243–1277.

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

**Malvey, Paul and Christine Archibald**, “Uniform Price Auctions: Update of the Treasury Experience,” 1998.

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

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

Table 1: Summary Statistics

Maturity	No. Auctions	Mean Issue Size (\$bn)	% Demand Tendered			% Won		
			Primary	Direct	Indirect	Primary	Direct	Indirect
CMBs <sup>a</sup>	101	24.6	88%	6%	6%	76%	7%	17%
4 week	222	31.9	85%	7%	8%	65%	8%	27%
13 week	222	28.5	86%	7%	7%	67%	8%	25%
26 week	222	24.9	84%	7%	9%	61%	8%	31%
52 week	55	24.3	82%	8%	10%	61%	10%	29%
2 year	51	34.7	73%	13%	14%	55%	17%	28%
5 year	51	34.3	70%	12%	18%	48%	12%	40%
10 year	51	20.5	69%	13%	18%	46%	16%	38%

<sup>a</sup> “CMBs” refers to the Cash Management Bills.

<sup>b</sup> “Primary” refers to the Primary Dealers, “Direct” to the Direct Bidders, “Indirect” to the Indirect Bidders.

Table 2: Description of Bids

Maturity	Bid <sup>a</sup>		Within Auction SD[Bid] <sup>a</sup>		% of Issue Size <sup>a</sup>		% of Tender Won <sup>a</sup>	
	Primary	Direct	Primary	Indirect	Primary	Direct	Primary	Indirect
CMBs <sup>b</sup>	0.1501	0.1389	0.0244	0.0201	19%	5%	21%	64%
4 week	0.0943	0.0699	0.0254	0.0337	18%	3%	19%	84%
13 week	0.1119	0.0866	0.0248	0.0332	19%	3%	19%	84%
26 week	0.165	0.1368	0.0275	0.0391	20%	4%	16%	71%
52 week	0.2617	0.2356	0.0299	0.0333	17%	4%	20%	67%
2 year	0.5604	0.5231	0.0397	0.046	13%	4%	22%	70%
5 year	1.5627	1.4902	0.0682	0.0631	10%	3%	24%	82%
10 year	2.7229	2.6482	0.0732	0.0706	11%	3%	21%	71%

<sup>a</sup> "Bid" refers to quantity-weighted bid-yield averaged across auctions. "Within-Auction SD[Bid]" refers to the within-auction standard deviation of quantity-weighted bid-yield, averaged across auctions. "% of Issue Size" refers to the percentage of total issue size demanded by a bidder, averaged across auctions. "% of Tender Won" refers to the percentage of the tender quantity that the bidder won at the auction, averaged across auctions.

<sup>b</sup> "CMBs" refers to the Cash Management Bills.

<sup>c</sup> "Primary" refers to the Primary Dealers, "Direct" to the Direct Bidders, "Indirect" to the Indirect Bidders.

Table 3: Description of Normalized Bids<sup>a</sup>

Maturity	Normalized Bid <sup>b</sup>			Within Auction SD[Norm Bid] <sup>b</sup>		
	Primary	Direct	Indirect	Primary	Direct	Indirect
4 week	2.04	-0.36	-2.73	2.53	3.37	2.66
13 week	1.99	-0.53	-2.36	2.48	3.32	2.49
26 week	2.33	-0.48	-16.3	2.75	3.91	2.72
52 week	2.68	0.07	-0.79	2.99	3.33	1.70
2 year	4.84	11.12	-2.36	3.97	4.60	9.39
5 year	8.01	0.77	-4.41	6.82	6.31	12.44
10 year	7.84	0.37	-5.39	7.32	7.06	19.20

<sup>a</sup> Normalized Bids are defined as Quantity-weighted bids (in basis points) minus the interest rate of the corresponding maturity reported by the U.S. Treasury on the day of the auction.

<sup>b</sup> “Normalized Bid” refers to quantity-weighted normalized bid-yield averaged across auctions. “Within-Auction SD[Norm Bid]” refers to the within-auction standard deviation of quantity-weighted normalized bid-yield, averaged across auctions.

<sup>c</sup> “Primary” refers to the Primary Dealers, “Direct” to the Direct Bidders, “Indirect” to the Indirect Bidders.

Table 4: Analysis of Bids

Dependent Variable	Bills		Notes	
	(1) QwBid(bp) <sup>a</sup>	(2) QwBid(bp)	(3) QwBid(bp)	(4) QwBid(bp)
Direct	-2.457 (0.0580)	-0.929 (0.0600)	-5.974 (0.270)	-0.965 (0.314)
Indirect	-4.204 (0.0604)	-2.529 (0.0613)	-10.89 (0.356)	-4.437 (0.399)
%Q Total <sup>b</sup>		10.04 (0.219)		61.75 (5.452)
Constant	13.87 (0.0316)	11.99 (0.0426)	172.0 (0.261)	165.0 (0.460)
Observations	41,359	41,359	13,692	13,692
R-squared	0.254	0.289	0.086	0.099
Number of auctions	822	822	153	153

<sup>a</sup> “QwBid(bp)” refers to the Quantity-weighted Bids, reported in basis points.

<sup>b</sup> %Q Total is the total (maximal) demand.

<sup>c</sup> Auction fixed effects are controlled for in every specification.

<sup>d</sup> Robust standard errors, clustered by auctions, are reported in the parentheses.



Table 5: Analysis of Bid Shading

Dependent Variable	Bills				Notes			
	(1) Shade 1 <sup>b,c</sup>	(2)	(3)	(4)	(5) Shade 1	(6)	(7)	(8) Shade 2
Direct	-1.904 (0.0851)	-1.434 (0.0969)	-0.862 (0.0727)	-0.771 (0.0884)	-4.696 (0.305)	-2.203 (0.255)	-0.0954 (0.0103)	-0.0480 (0.0105)
Indirect	-3.511 (0.105)	-2.996 (0.117)	-1.125 (0.0813)	-1.025 (0.0978)	-13.36 (0.684)	-10.15 (0.469)	-0.122 (0.0116)	-0.0608 (0.0129)
%Q Total <sup>e</sup>		3.085 (0.353)		0.600 (0.330)		30.73 (4.399)		0.584 (0.108)
Constant	0.841 (0.0543)	0.265 (0.0819)	1.174 (0.0441)	1.062 (0.0756)	-1.888 (0.478)	-5.394 (0.353)	0.125 (0.00883)	0.0579 (0.0122)
Observations	41,264	41,264	41,264	41,264	13,692	13,692	13,692	13,692
R-squared	0.095	0.097	0.015	0.015	0.158	0.162	0.062	0.069
Number of auctions	822	822	822	822	153	153	153	153

<sup>a</sup> %Q Total is the total (maximal) demand.

<sup>b</sup> Shade 1 ( $\equiv B(\theta)$ ) is as in Definition 1 and Shade 2 ( $\equiv S(\theta)$ ) is as in Definition 2.

<sup>c</sup> Bid shading is reported in basis points.

<sup>d</sup> Auction fixed effects are controlled for in every specification.

<sup>e</sup> Robust standard errors, clustered by auctions, are reported in the parentheses.

Table 6: Analysis of Bid Shading at the Marginal Step

Dependent Variable	Bills				Notes			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Shade 1 <sup>b,c</sup>		Shade 2		Shade 1		Shade 2	
Direct	-1.237 (0.0459)	-1.085 (0.0446)	-1.940 (0.577)	-1.841 (0.570)	-2.428 (0.167)	-2.345 (0.165)	-0.0185 (0.0019)	-0.0141 (0.0015)
Indirect	-2.192 (0.0497)	-2.049 (0.0.0456)	-2.619 (0.589)	-2.526 (0.0.584)	-5.844 (0.217)	-5.661 (0.191)	-0.022 (0.002)	-0.012 (0.0013)
%Q Marginal <sup>a</sup>		6.214 (0.554)		4.041 (4.333)		8.960 (2.468)		0.475 (0.0659)
Constant	-0.217 (0.0225)	-0.474 (0.0240)	2.693 (0.312)	2.526 (0.336)	-0.241 (0.136)	-0.485 (0.110)	0.0218 (0.00146)	0.0089 (0.00113)
Observations	36,860	36,860	36,860	36,860	12,034	12,034	12,034	12,034
R-squared	0.161	0.169	0.002	0.002	0.047	0.047	0.044	0.088
Number of auctions	822	822	822	822	153	153	153	153

<sup>a</sup> %Q Marginal is the total demand at the marginal step.

<sup>b</sup> Shade 1 ( $\equiv B(\theta)$ ) is as in Definition 1 and Shade 2 ( $\equiv S(\theta)$ ) is as in Definition 2.

<sup>c</sup> Bid shading is reported in basis points.

<sup>d</sup> Auction fixed effects are controlled for in every specification.

<sup>e</sup> Robust standard errors, clustered by auctions, are reported in the parentheses.

Table 7: Bidder Surpluses: July 2009-October 2013

	Primary		Direct		Indirect	
	(bp)	(M\$)	(bp)	(M\$)	(bp)	(M\$)
CMBs <sup>a</sup>	0.17	40.6	0.02	3.8	0.04	9.6
4-Week	0.04	26.2	0.00	2.1	0.002	1.1
13-Week	0.13	86.4	0.02	11.1	0.008	5.3
26-Week	0.33	183	0.03	15.1	0.026	14.6
52-Week	0.68	90.8	0.08	10.5	0.14	18.4
2-Year	7.40	1310	1.15	202	0.91	161
5-Year	13.07	2280	1.87	326	1.39	243
10-Year	22.22	2320	3.58	373	1.73	180
Overall	2.3	6337	0.35	943.5	0.23	633

<sup>a</sup> “CMBs” refers to the Cash Management Bills. “(bp)” refers to the basis points, “(M\$)” to the million U.S. dollars.

<sup>b</sup> “Primary” refers to the Primary Dealers, “Direct” to the Direct Bidders, “Indirect” to the Indirect Bidders.

Table 8: Analysis of Bidder Surpluses

Dependent Variable	Bills				Notes	
	(1) Surplus	(2) Surplus	(3) Surplus	(5) Surplus	(6) Surplus	(7) Surplus
Direct	-18.67 (1.662)	-12.14 (1.725)	-10.61 (1.566)	-1,466 (126.9)	-792.6 (89.68)	-443.4 (74.97)
Indirect	-25.08 (2.424)	-17.92 (2.403)	-15.61 (2.034)	-1,970 (162.3)	-1,101 (115.6)	-697.9 (97.91)
%Q Total <sup>b</sup>		42.90 (8.635)	20.29 (9.028)		8,305 (1,112)	7,087 (1,024)
# Indirects Observed			6.615 (1.609)			201.4 (27.60)
Constant	25.73 (1.237)	17.71 (1.678)	16.07 (1.586)	2,007 (121.6)	1,059 (96.24)	669.9 (95.08)
Observations	41,264	41,264	41,264	13,692	13,692	13,692
R-squared	0.032	0.034	0.041	0.329	0.359	0.392
Number of auctions	822	822	822	153	153	153

<sup>a</sup> Surpluses are in \$1,000s.

<sup>b</sup> %Q Total is the total (maximal) demand.

<sup>c</sup> Auction fixed effects are controlled for in every specification.

<sup>d</sup> Robust standard errors, clustered by auctions, are reported in the parentheses.

Table 9: Allocative Efficiency of the Auctions

Maturity	Efficiency Loss <sup>a</sup> (in bp)
1-month	0.67
3-months	0.68
6-months	0.76
12-months	0.65
2-year	2.08
5-year	4.50
10-year	6.41

<sup>a</sup> Efficiency Loss is the amount of surplus lost due to misallocating the notes/bills.




Figure 1: Illustration of Bid Shading when Residual Supply is Known