QUANTITY COMMITMENTS IN MULTIUNIT AUCTIONS: EVIDENCE FROM CREDIT EVENT AUCTIONS

ERIC RICHERT

Abstract. Credit Default Swaps (CDS) are financial derivative products that insure bond investors against firm-default. Determining the payout, however, is complicated because the outstanding value of the insurance is larger than the debt outstanding. The value of a bond is also heterogeneous. CDS payouts are therefore determined in a two-stage auction. In the first stage dealers commit to either supply or purchase a fixed quantity at the unknown final price. Then, the excess supply or demand is announced and a multiunit uniform price auction is held to determine the market clearing price. Dealers have an incentive to bid strategically; in addition to the standard information rents in multiunit auctions, the two stage auction features (i) learning across rounds, (ii) pre-committment of quantities in the first round, and (iii) heterogeneous positions in CDS contracts. The paper develops and estimates a structural model of bidding behavior in these auctions and uses it to quantify the role of each of these channels in the dynamic auction process. I then consider counterfactual changes to the auction format, including a double auction design with step function bidding.

Acknowledgements: PRELIMINARY AND INCOMPLETE. I thank my supervisors Jason Allen, Robert Clark, and Brent Hickman. Helpful comments were provided by Ali Hortaçsu, Kim Huynh, Jakub Kastl, Marzena Rostek, James Thompson, Andreas Uthemann, Milena Wittwer, Yu Zhu and from seminar participants at The Bank of Canada, Tilburg, Chicago, Wisconsin, HEC Montreal, Western, Kellogg, Johns Hopkins, McGill, NYU, Queen's, Washington University in St. Louis, Analysis Group, Charles River Associates, the CEA meetings and the IO Finance Reading Group at the University of Wisconsin. I acknowledge support from the Social Sciences and Humanities Research Council of Canada. Correspondence to: Eric Richert - Princeton University; Email: richert.eric@gmail.com.

Date: November 14, 2022.

1. Introduction

U.S. firms issue nearly $2 trillion in corporate debt per year. The large institutional investors who purchase this debt will often hedge against default risk by buying insurance using Credit Default Swaps (CDS).[1] CDS are derivative contracts which provide insurance against a credit event (eg. bankruptcy, failure-to-pay, restructuring) occurring on some set of bonds. CDS contracts, however, are not only used for hedging, but also for speculation—investors betting on firm default. This insurance market has a gross notional volume outstanding of around $10 trillion and credit events have included Fannie Mae, Lehman, Greece and GM.[2]

The target payment for CDS contracts is the difference between the par value on a bond and its recovery value. The focus of this paper is in the determination of recovery value. When CDS contracts were first introduced, settlement involved a physical transfer of bonds from buyers of insurance to sellers, in exchange for the cash value of the insurance. This arrangement, however, is complicated by the fact that the market can have many more CDS contracts than bonds. Buyers of insurance do not necessarily own the bond to physically settle. Physical settlement then would result in a short squeeze where the few investors who own the bond can charge a very high price to investors needing to source it in order to realize the insurance payout. As a result, the participants in this market and the International Swaps and Derivatives Association (ISDA) agreed to instead settle these contracts using a cash payment, the value of which is determined by holding a two-stage auction for bonds. However, because investors are active in both the auction and hold the underlying contracts they may have incentives to distort auction prices to increase contract payoffs.[3]

The effectiveness of auction-based cash settlement depends on the ability of bidders to influence the final auction price. This is a concern because the CDS market is made up of only a few large players who could be in a position to exert their market power.[4] These same players had a central role in the transition to cash settlement and the set up of the auction mechanism.[5] In this paper I estimate a structural model of bidding in the current environment and use the model estimates to quantify the distortions arising from market power. I then consider alternative auction formats, which I argue can increase market efficiency. I show that the current design results in prices that are influenced by dealers' insurance positions. Since in most (but not all) auctions dealers are net holders of insurance, this leads to a price that is on average 4 cents on the dollar below the fair insurance amount. Given the size of the market in an auction such as the one following

[1]https://www.sifma.org/resources/research/us-corporate-bonds-statistics/

[2]https://www.isda.org/a/JUPTE/Global-CDS-Market-Study.pdf

[3]A similar incentive is studied in a theoretical model of spot markets for derivatives, by Zhang (2022).

[4]Some non-dealer firms have expressed a desire for more direct participation in the auctions, see Rutledge (2009).

[5]A recent lawsuit New Mexico State Investment Council v. Bank of America et al. case number 1:21-cv-00606 in the U.S. District Court for the District of New Mexico, alleges that the large dealers used their market power to influence clients acceptance of the auction process and were heavily involved in its design.

GM's default, this would cost sellers of insurance an additional $2.6B. Because the size and direction of the bias in any given credit event depends on the private net insurance positions of dealers, investors cannot offset the bias by adjusting the amount of insurance purchased. While investors may adjust their insurance positions to account for the average bias (across credit events), doing so leaves them bearing additional risk. By changing the auction rules I show the price bias can be reduced by about 41%. Investors are also left holding risk from the variance in auction outcomes that arises due to dealers' price impacts. I show that under alternative auction rules the standard deviation of auction outcome risk can be cut by 30%. The decreased ability to hedge risk for investors may mean that investors require additional compensation to own bonds, which has important implications for the costs of funds to firms/governments.

At each auction, bonds are bought and sold by large investment banks (Barclays, Goldman Sachs, etc.) to determine their value.[6] At the auction, bidders pay or receive the auction price for the bonds they buy/sell, and they receive one minus the auction price to cash-settle their CDS contracts. The current auction format for settling CDS contracts involves a nonstandard two-stage design. In a first stage, the auctioneer accepts initial quantity commitments, i.e. buy and sell commitments that are enforced in the second stage. This first stage determines whether there is an excess demand for bonds or an excess supply. In the second stage the auctioneer uses a multi-unit auction to determine the price which clears the market.

To fix ideas, let us consider an example of the behavior of a bidder (for instance Barclays) following a credit event. At the start of the auction suppose that Barclays' position is $4M of bonds and $10M of insurance (CDS contracts). Barclays would like to receive a low price for their bonds in order to maximize their CDS payoff. In the first stage, because they are a buyer of insurance, they can commit to supply bonds.[7] This commitment is costly if Barclays has a high value for the bond — by making a first stage commitment of $3M in bonds, Barclays reduces their exposure to the auction price since it lowers their effective CDS position to $7M.[8] This has the benefit for the auctioneer of decreasing bidders' incentives to shade their bids in the second stage. After the first stage, the auctioneer sums the orders across participants and publicly announces whether the second stage is for excess supply or excess demand and the size of this excess. This announcement reduces the uncertainty each bidder faces about the degree of competition. If the second stage game allowed for both supply and demand bids following the initial round, the initial round quotes would be non-binding and non-informative. However, the single directional second round means that sometimes a bidder will be unable to adjust the change in position from their initial quote and this provides a cost for a particular choice in the first round. In addition to their first-stage quantity commitments, all bidders are required to

[6]After the auction the bidder can hold the bonds or sell them to their clients in the secondary market.
[7]If they had instead been a seller of CDS they would have been limited to buying in this stage.
[8]Barclays is paid one minus the auction price, for all their $10M of CDS contracts, they receive $3M times the auction price for the bonds sold and so have a final exposure of $7M. This reduction in exposure assumes that the post-auction price is not driven by the auction price. Results in Table 7 show this is supported by the data.

submit price quotes. The average quote can be thought of as the common value for the bond. It does not bind bidders in the second stage, but helps aggregate individual bank signals about the value of the bond.

The cases of excess supply and excess demand after the first stage must be considered separately. If there is excess demand, only supply bids are allowed. This means Barclays is excluded from expressing a desire to buy bonds in the auction, thereby reducing competition.[9] If instead the first stage results in excess supply, Barclays will bid below its value in order to (i) extract information rents, and (ii) increase the amount it is owed on its CDS position (as a buyer of CDS, Barclays wants a lower auction price to increase the amount of insurance they are owed). Consistent with the existence of strategic bidding, bond prices in the auctions are usually a few cents on the dollar below their secondary market counterparts, see, for example, (Coudert and Gex. (2010), and Gupta and Sundaram (2012)). Given the large size of the CDS market, a few-cent shift can lead to a hundred million dollar change in the amount due from CDS sellers. The existence of this distortion makes the auction design crucial for valuation of CDS contracts and the functioning of the insurance market.[10] For example, the international accounting standards board lists the effect of the auction process as an important reason why CDS prices may not be the best measure of the inherent credit risk. The estimates from my model allow me to quantify this distortion without relying on the post-auction transaction prices, which only exist for one third of the auctions, as proxies for bidder values.

I evaluate if a counterfactual change to a double auction, where bidders simultaneously submit supply and demand orders, can reduce these distortions.[11] The distortions from participation constraints have been illustrated in theoretical work on these auctions, see Du and Zhu (2017). They demonstrate that a move to a double auction would reduce the price bias. However, their result relies on assumptions that bidders have zero-average CDS positions and that they have common rates of decreasing marginal values. These assumptions are rejected by the data. Because dealers' are estimated to be mostly net buyers of CDS, the incentives for buyers and sellers to distort their bids no longer cancel out in the double auction design and so the ranking of the double auction and current format must be determined empirically. In addition, because I allow dealers to account for their price impact, learning and first stage position reductions may benefit the auctioneer in the current format. Relative to the current format, the double auction structure eliminates the reduction in exposure from position reductions, increases the uncertainty about opponents' demands, but increases participation in the auction by removing constraints on the direction of eligible bids.

To evaluate the effect of hypothetical changes to the auction rules, I compute the equilibrium strategies of bidders under counterfactual auction scenarios. To compute these

[9]In the first stage, Barclays could not sell bonds because they were a net buyer of CDS, while in the second stage only demand orders were accepted.

[10]The functioning of the insurance market also impacts the bond market, for example Das et al. (2014)

[11]In settings with imperfect competition the double auction is not fully efficient because bidder's strategically account for the price impact of their bids. This is true in the models of Kyle (1989),Vives (2011), and Ausubel et al. (2014).

strategies I apply a method which I develop in a companion paper Richert (2021). Direct computation of equilibrium in multi-unit auction settings has been elusive. This has limited the counterfactuals considered to exercises that provide an upper bound on the benefits of eliminating bid shading (e.g. Hortaçsu and McAdams (2010), Kastl (2011)).[12] The main challenge with the equilibrium computation is that bidders' strategy functions are high dimensional and complex. This means that both Euler-based approaches (which solve the strategies by taking sequences of steps along a path described by a differential equation), and approaches based on parametrizations of the strategy functions (e.g. Armantier et al. (2008)), cannot be applied. Instead, my method begins by guessing a data-generating process (DGP) for equilibrium bids and adjusts the DGP until the distribution of values that rationalizes those bids given the rules of the game matches the true distribution of values (the known structural primitive when solving the counterfactual). This approach does not require imposing parametric restrictions on the bid strategy functions, insures that the equilibrium constraints are satisfied exactly at the solution and in a single execution can solve for the entire set of counterfactuals consistent with estimates from a set-identified model.

Unlike in standard multiunit auction settings where the econometrician is only interested in learning a bidder's private value from their bids, in credit event auctions we need to identify both bidder private values and their CDS positions. I extend identification arguments from multi-unit auctions by using restrictions on the shape of the marginal value curve (ie. bounded and weakly decreasing in quantities) to jointly bound the set of CDS positions and marginal values for every bidder. It is important to separate these two components because bidders use their bids not just to express their value for the bonds but also to influence the payments due on their CDS positions. Recalling the Barclays example, when a CDS buyer raises its bid, it increases the expected auction price, which reduces the payments it will receive on its CDS position. This effect increases the incentive to shade bids for CDS buyers and reduces shading of CDS sellers compared to the shading due only to private information; for an overview see Kastl (2020).

Model estimates allow me to document the size of the distortions due to information rents. Because dealers often own insurance, they have additional incentive not to bid aggressively in the auction and therefore auction prices are often below bond prices in the secondary market, resulting in larger insurance payments. The estimated price bias is similar to the gap between the auction price and secondary market prices of bonds. I use these estimates to quantify several different strategic channels described in theoretical work on these auctions. The source of these frictions is important for policy: theoretical work focusing on different channels have proposed different solutions. Chernov et al. (2013) propose pro-rata rationing; Du and Zhu (2017) propose a double auction; and Peivandi (2015) suggests a fixed price.[13] I perform several tests that suggest the first

[12]These comparisons are informative in the Treasury setting because the degree of bid shading is small. In the credit event auction context these bounds suggest considerable improvement may be possible.

[13]Chernov et al. (2013) use an environment featuring perfect information and common values for the defaulted bond to highlight the role of short sale constraints (difficult to short sell the bonds) and constraints prohibiting some participants from holding defaulted bonds. Rather than have pro-rata

round price quote captures most of the common value element in these auctions. Given this I focus on the double auction policy proposal from Du and Zhu (2017) to reduce market power in assuming independent private values (IPV).

To understand the strategic channels in the current settlement format I perform a decomposition exercise which allows me to separately quantify the roles of (i) the reduction in asymmetries due to first-round quantity commitments, (ii) the reduction in uncertainty from learning the excess supply/demand available (iii) the information learned about opponents' values from the endogeneity of the excess supply/demand. The first stage commitments reduce bidders' exposure to the final price and should reduce their incentive to shade bids in the second stage. However, results from the decomposition suggest that the position reduction from the first stage commitments do not result in large changes in bids. This is driven by the fact that dealers submit only part of their position and roughly 60% of bidders make zero commitment in the first stage. The quantity announcement reduces the aggregate uncertainty and because the quantity is a function of the opponents' choices, provides information on the level of competition. The decomposition results show that the announcement of the open interest prior to bidding has a net pro-competitive effect, mainly through its role in reducing the total level of uncertainty that bidders face. Unfortunately this decomposition does not allow us to separate the role of participation constraints because the main effect operates by making the first stage commitments bind. Relaxing these constraints would therefore lead to different first stage choices, which cannot be captured in this exercise.[14]

Relative to the current auction format, the double auction design eliminates the effect of learning and the role of participation constraints, and increases the effective CDS positions of dealers at the auction. Given that the estimated positions of the dealers are more often as net buyers of CDS, they have a common incentive to bid less aggressively in these auctions. These estimated shading effects are large, for example the average price in the current format is 4 cents on the dollar below the price the bidders would be willing to pay for the bond.In the counterfactual exercise I show that the double auction provides an improvement in the functioning, decreasing the total gap from the competitive price by 41.25%.

---

rationing at the margin, they follow Kremer and Nyborg (2004) and propose pro-rata rationing. The idea is to eliminate under-pricing in the uniform price auction. Peivandi (2015) instead focuses on the problem of bilateral settlement before the auction and uses a mechanism design approach to show that the optimal mechanism to control this problem is to use a fixed price which is independent of signals. Finally, Du and Zhu (2017) focus on the role of the constraints on behavior across rounds and the inefficiencies generated by the excluded participants. They model behavior in an IPV setting where the dealers have no price impact and hold zero average CDS positions and show in that context a double auction is efficient and would improve performance.

[14]If you calculate only the change in bids at the second stage, holding fixed the distribution of open interests and opposing bids this only allows the constrained bidders to adjust their response. However, in equilibrium the other bidders would respond to this adjustment and open interests would no longer play a role in second stage bidding.

Section 2 presents details of the CDS auction institution and introduces the data, 3 introduces the model, 4 discusses identification, 5 presents the estimation 6 considers the results and section 7 examines the counterfactual experiments.

2. Institutions and Data

CDS are financial derivatives which provide insurance against a pre-determined set of credit events (eg. bankruptcy) occurring on a pre-specified set of bonds.[15] These contracts initially used *physical settlement*, akin to basing settlement on scrapage value in other insurance markets. In *physical settlement*, the insurance buyer delivers the bond to the insurance seller and in return receives the par value of the bond. This leaves the buyer with the full initial value, and the insurer can claim any bond recoveries. However, in addition to bond owners, CDS contracts may be purchased by speculators that do not own the underlying bond. These so-called *Naked* CDS contracts where the buyer does not not own the underlying bond allow speculators to use CDS to bet on the creditworthiness of the company.[16] The presence of speculators means that the volume of CDS is often many times the outstanding volume of bonds, and so physical settlement of all contracts would require the bonds to be recycled through the market. This could produce a short squeeze in the bond market, preventing physical settlement from providing fair insurance for naked buyers (Gupta and Sundaram (2015)). Physical settlement also produces an inefficient allocation of bonds when some CDS buyers have a higher value than the sellers of holding the bond through the recovery process.

These issues were anticipated in the lead up to the default of Delphi in 2005, where there was $25B net notional of CDS contracts written on $2B of bonds. To address the problems linked with physical settlement, the twelve big dealer-banks decided to settle contracts in cash at a price determined in an auction for the underlying bonds. A two-stage auction design was proposed to allow participants to replicate the outcomes of physical settlement. In the first stage dealers submit physical settlement requests and in the second a uniform price multi-unit auction is held to clear the market. To replicate physical settlement, dealers could submit requests to buy/sell in a first stage for as many bonds as they would have transferred physically, resulting in the same set of payments and transfers. Following 2009 these auctions were written into all CDS contracts as the settlement mechanism.[17]

[15]For a survey of the literature on CDS markets, contract terms, and pricing, see Augustin et al. (2014).
[16]Since the market for CDS on a particular company is thin and illiquid it has thereby been argued that allowing this type of contract is efficient as it helps create liquidity and information, reducing market frictions.
[17]There have been two major changes in this market since the first auctions in 2006: the big bang and small bang protocol. The main effect of these rules (effective 2009) were to tie the CDS contract payouts to the auction prices. Credit events can be on bankruptcy, failure-to-pay or restructuring from the underlying obligations. The big bang protocol hardwired the auction process for bankruptcy and failure-to-pay while the small bang did the same for restructuring events.

There were 209 credit events between 2006 and the Fall of 2019. Of these, 84 were loan credit default swaps (LCDS) and 125 CDS.[18] There were 7 auctions which did not proceed to the second stage, and in 16 cases auctions were not held, resulting in a sample of 186 auctions.[19] I collect data from *creditfixings.com*, (administered by Creditex and MARKIT) on all bids made in credit event auctions, and obtain the lists of eligible bonds from the determinations committee.[20] The covered entities included companies and countries; eligible debt includes corporate and sovereign bonds, syndicated loans, commercial mortgage backed securities (CMBS) and mortgage backed securities (MBS).[21] Summary statistics are displayed in Table 1. The auctions have an average of 11 participants, usually the nine global dealers and two largest regional participants. The price determined in the auction averages 43.44 cents to the dollar. That is, the auction price for the bond is 43.44 percent of the par value of the bond; there is, however, substantial variation across auctions. The IMM, which is an average of price quotes given by dealers at the first stage of the auction, is fairly similar to this final price. NOI (net open interest) denotes the volume of excess supply or demand for bonds resulting from the first stage. In total, 355 bidders submit requests to buy in the first stage of one of the 186 auctions, 535 submit requests to sell, and 1167 submit zero quantity bids in the first stage. Almost 69 percent of the auctions result in excess supply in the first stage.

From the determination that a credit event has occurred to the final payout, the auction process is administered jointly by a committee whose members are determined based on their global notional volumes. These large dealers are obliged to participate in most auctions: failure to participate could threaten their eligibility to participate in future auctions. The committee also determines the final set of deliverable obligations for the auction. In cases where the issuer has debt of multiple maturities or risk levels the auction may be held separately on different buckets. In these cases bonds eligible for submission in the shorter maturity or higher security level can also be delivered into the lower auction. This means that the bonds that can be delivered in a particular auction are systematically homogeneous. Further, bidders will not submit bonds at random but will first submit the cheapest-to-deliver bonds. The total volume of bonds at auction is only a fraction of all eligible bonds. For each eligible bond, I obtain volume and trait information

[18]LCDS are similar to CDS contracts but have loans rather than bonds as the underlying reference obligation.

[19]In the design of the auctions the ISDA determined a set of situations where an auction is not required to be held. This occurs if for certain maturity buckets there are no deliverable obligations in the bucket that is not shared with a shorter-dated bucket or if the determinations committee decides an auction on that bucket is not warranted due to limited notional volume of transactions within the bucket. In the first of these cases the price can be set by rounding down to the previous, shorter dated bucket. However, if at least 300 transactions are triggered after the restructuring credit event determination in the given maturity bucket and at least five dealers are parties to these transactions an auction must be held. The auctions which did not proceed to the second stage had no excess supply/demand in the first stage to be sold in the second stage.

[20]Details are provided at https://www.cdsdeterminationscommittees.org/credit-default-swaps-archive/

[21]Results are similar if sovereign bonds are excluded from the analysis. These may have different information environments, and the participant banks may hold these for different reasons. Sovereign bonds are only the underlying obligation for 5 auctions in the sample.

TABLE 1. Auction Description

The following table presents summary statistics for the auctions. Price is the final market clearing price from the auction. IMM is the initial market midpoint calculated using bidders' first round price quotes. The NOI is the excess supply or demand from summing over each bidder's quantity commitments. Probability to buy takes a value 1 if the auction results in excess supply (accepts demand bids at stage two).

	N	Mean	Sd	P10	p50	P90
N Dealers	186	11.09	2.51	8	11	14
Price ($.01)	186	43.34	32.66	4.06	35.56	91.20
IMM	186	43.44	31.93	4.76	34.50	88.66
$\|NOI\|$ ($millions)	186	94.64	167.53	2	36.53	233.46
NOI % Bonds	186	16.67	1.03	-12.20	0.27	25.73
Probability to buy	186	0.688				

from Bloomberg and obtain loan information from DealScan. Summary statistics for the deliverable obligations in each auction are provided in Table 2. For 56 of the auctions, the bonds are covered by TRACE and so I also obtain trading data for the bond around the auction date.[22] On average, 11 bonds are eligible for submission into the auction. The bond characteristics vary substantially across auctions.

TABLE 2. Some Other Bond Descriptives

For each auction this table summarizes various features of the set of eligible bonds. In total there are 2,004 eligible bonds across all auctions. FRN % denotes the share of the eligible bonds that are floating rate notes (coupon payments linked to a benchmark rate, usually LIBOR).

Variable	Obs	Mean	Std. Dev.	Min	Max
Number of bonds	186	10.774	27.699	1	298
Max maturity (years)	186	10.419	11.928	0	50
Min maturity (years)	186	2.968	4.913	0	50
Max coupon %	186	5.5	4.8	0	29.5
Min coupon %	186	2.7	3.3	0	11.8
Share FRN %	186	33.8	41.7	0	1

The within-auction variation is summarized in Table 3. The table shows summary traits of the bonds, including volume, duration, convexity and conversion factor. Duration and convexity measure the exposure of the bond to interest rate risk. Duration has substantial variability within auction which is heavily influenced by the fact that the set of eligible bonds often contains some share of floating rate notes (FRN) and some share of long-term coupon bonds. Convexity is much more similar within auctions than across. The volume of individual issues varies substantially within auction. Because the dealers should anticipate which bonds are cheapest to deliver, they should have common expectations

[22]The Trade Reporting and Compliance Engine is the FINRA database for the mandatory reporting of over-the-counter transactions in eligible fixed income securities. Broker-dealers have an obligation to report transactions in eligible securities under an SEC-approved set of rules.

about the set of bonds in this pool, so the estimated values will represent bidders' values for that set of bonds.[23]

TABLE 3. Bond Measures

Volume ($B), duration convexity and conversion factor as calculated for each bond in the eligible set that can be submitted to the auction. Each variable x_{ja} for auction j admissible bond a has between variable x_j and within $x_{ja} - x_j + x$, where x is the global mean. While the "within" reported minimum eg. for volume is negative, this does not indicate negative volume of any issuance but refers to the deviation from each auctions average issuance size and naturally, some of those deviations must be negative. Across 186 auctions there are a total of 2,004 eligible bonds.

Variable		Mean	Std. Dev.	Min	Max
volume	overall	7.21	85.63	0	3500
	between		9.52	0	126.79
	within		81.44	-108.15	3380.41
duration	overall	3.33	10.85	0.86	127
	between		9.77	0.87	103
	within		9.50	-31.34	123.62
convexity	overall	0.86	1.49	0	10
	between		1.33	0	10
	within		1.21	-3.62	10.38

CDS contracts are traded over-the-counter and disaggregated trading data are not available.[24] Prior to 2010 reporting requirements for these transactions were limited. After 2010, information on all standardized and confirmed CDS transactions involving U.S. entities was reported to DTCC.[25] This data is available to regulators through the DTCC's Trade Information Warehouse. Paulos et al. (2019) uses the regulatory filings from the DTCC from 2014-2017 for the subset of dealers regulated by the Federal Reserve. They show that in the auction most of the dealers are net buyers of protection. The positions data that are sold to market participants do not contain enough information to reconstruct dealers net open positions on individual entities. Further, it is unlikely that dealers can reconstruct precise estimates of the other dealers' positions as they often transact with each other through inter-dealer brokers to preserve anonymity. Another indication of the

[23]If counterfactual changes in the auction format resulted in a large change in the total volume of bonds exchanged, one may then worry that the composition of the pool (and hence bidders' values) may vary, however in all exercises presented here the volume of bonds is comparable across the formats.

[24]Bidder's inventories influence their bidding behavior in CDS auctions. Therefore other bidders, and the econometrician are interested in these levels. However, with the available data these inventories are not observed and tools to construct inventories from observed trades such as Hansch et al. (1998) cannot be used.

[25]Paddrik and Tompaidis (2019) use this data to examine the costs dealers face to act as market makers.

lack of information on others' positions is that netting does not occur in this market. If an opponent's position was precisely known, dealers would likely engage in netting to free up collateral and reduce exposure to counter-party risk. Following the financial crisis, some CDS have moved to central clearing but this tends to be index CDS and more liquid companies, see for example Slive et al. (2012), and has not affected the single names on which credit events occurred.[26]

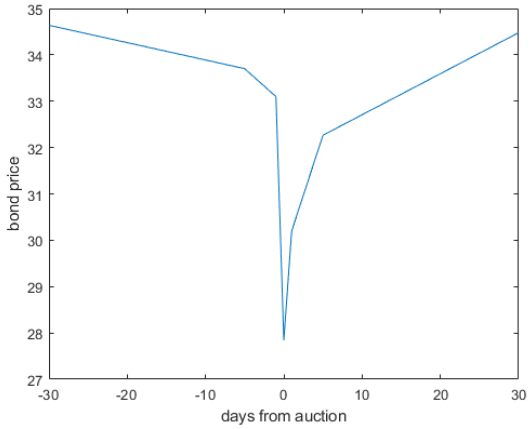
2.1. **Evidence of market power.** Figure 1 plots the average transaction prices in the secondary market (and the auction price on day 0). The V-shaped pricing pattern is consistent with the findings in existing papers that analyze this difference. Coudert and Gex. (2010), and Gupta and Sundaram (2012), for example, document a large gap between a bond's price on the auction date and secondary market prices around the auction day, showing that the auction price tends to be below both the pre- and post-auction trading price.

Although this price gap is consistent with the presence of market power it could also come from other sources. For example, dealers may take on larger bond positions around the auction, with clients selling both CDS and bonds to dealers and so the prices around the auction may reflect a larger discount for the additional inventory risk. There could also be additional risk in the bond price around the auction as auction outcomes may reveal information about the bond value to bidders. This relationship is discussed in Table 7. While the auction price itself has no explanatory power for post-auction bond prices,

[26]Single name CDS has a reference obligation or bond issued by a single issuer. Index CDS are credit securities on a basket of credit entities.

FIGURE 1. Average Secondary Market Bond Prices

The x-axis plots the number of days from the auction. The y-axis plots the average bond price per dollar.



the post-auction prices are correlated with the IMM, suggesting that some information relevant for the secondary market may be revealed during the first-stage of the auction.[27]

[27]In addition, the concentration of the quantity of bonds won by the largest winner does not explain the price gap around the auction date, suggesting the gap may not be due to differences in the expectations of dealers' actions through the recovery process.

2.2. **Current auction format.** The auction begins with a stage where bidders (i) submit initial quantities that they want to commit to settle at the final auction price and (ii) price quotes with a quantity and maximum spread that are pre-determined by the determination committee and are chosen depending on the liquidity of the defaulted assets. Following the first stage the auctioneer adds up all the quantity commitments and announces this along with the average price quote. They then hold a uniform price multiunit auction to clear the excess supply or demand. I illustrate the process with an example auction.

2.2.1. *Initial Quantity.* For an example of the initial quantities see Table 4. Quantities to buy and sell are summed across dealers to determine the Net Open Interest (NOI). In the example presented in Table 4 there is a NOI of \$47.397M. The auction is therefore in excess supply. The bidders are required to submit initial quantity submissions that are in the same direction as their net position in CDS. Therefore a bank that owns more CDS contracts than it has sold can only submit requests to sell. This direction restriction is set to replicate the transfers from physical settlement. This particular example auction had \$507M of eligible bonds.

To understand the incentives of dealers in this stage let us consider a dealer that is a net buyer of insurance. This dealer wants the auction to establish a low price for the bonds, which will result in a larger payout on their CDS position. However, by supplying additional units in the first stage, the dealer adjusts their final exposure to the auction price. On CDS which they own the dealer receives a cash settlement of $(1 - p^{auc})$ while they receive p^{auc} for bonds sold at the auction. Finally, the dealer must consider what opponents will learn from any realization of the NOI.

Figure 2 looks at the impact of the total quantity submitted by opponents in the first stage and the expected price in the auctions. This highlights the intuitive relationship

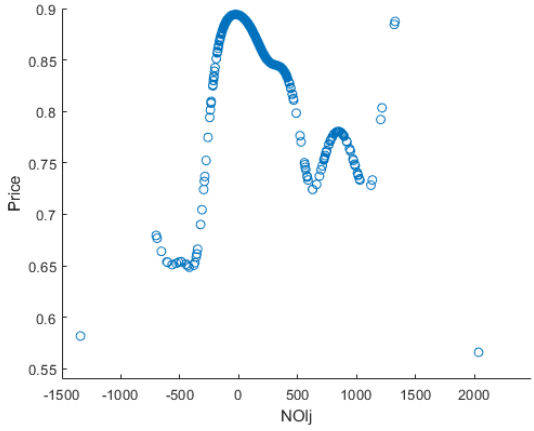
TABLE 4. Initial Round Quantities

This table presents the initial round quantities from the auction for Parker Drilling Co. An offer is a commitment to supply bonds. A bid is a commitment to buy bonds.

ID	Dealer	Bid/offer	Size (\$M)
1	Barclays Bank PLC	Offer	10
2	BNP Paribas SA	Offer	0
3	Credit Suisse	Offer	7.953
4	Deutsche Bank	Offer	0
5	Goldman Sachs International	Bid	6.53
6	J.P. Morgan Securities LLC.	Offer	26.974
7	Merrill Lynch, Pierce, Fenner & Smith Inc.	Offer	0
8	Morgan Stanley & Co. LLC	Offer	9.0
9	Societe Generale	Offer	0
	Subtotal Buying	Bid	6.53
	Subtotal Selling	Offer	53.927
	Total for Auction (Net Open Interest)	Offer	47.397

FIGURE 2. Expected Price Quantity Others

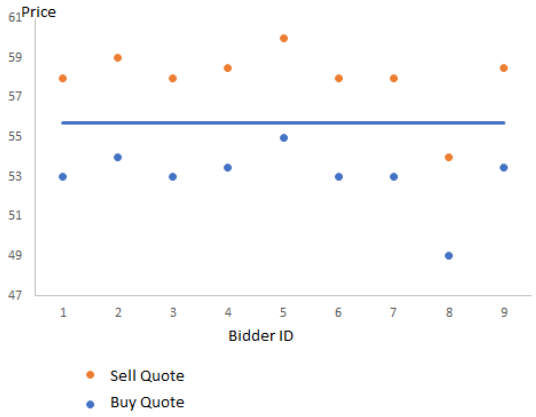
Nonparameteric smoothed estimates of the expected price as the $NOI_i = NOI -$(dealer j's commitment) varies. Expected price calculated as a fraction of the price cap, and expectations are taken by simulating residual supply curves which imposes the assumption that bids are conditionally independent in the second stage given the NOI submission of opponents and the bidders own submission.



that a small quantity to be cleared results in more competitive bidding, leading to a high price. However, as the quantity that needs to be cleared increases, in the second stage bidders can shade their bids more, and the expected price falls.

2.2.2. *Initial Quotes.* The first stage also includes a simultaneous submission of price quotes. The quotes are used to set a price floor (ceiling) in the auction when bidders are selling (buying) and are carried forward into the auction as part of the second stage bid. Prior to the second stage, the average quote is also announced. An example of the initial quotes in one auction is provided in Figure 3. The price caps serve to set a limit such that a dealer with a net CDS position larger than the total quantity being auctioned does not have an incentive to push the price to 100 or zero. To calculate this cap from the quotes, the auctioneer discards crossing/touching markets (where a buy price is above a sell price, and takes the 'best half': the highest half of the remaining bids and lowest half of the offers and calculates the average— called the Initial Market Midpoint (IMM). The spread is added to the average if it is an auction to buy or subtracted if it is to sell to set the cap. These quotes are carried over into the auction in the direction matching the NOI, at an auction-specific quantity, set in advance by the auctioneer. Any bids that are off-market are carried over at the IMM. Finally, the bids are used to determine fines for off-market bids. If an offer to buy is above an offer to sell, this indicates a trade-able market, and the off-market party will be fined based on the size of this difference multiplied by the fixed quote size. The maximum spread, as well as the quantity which is used to determine fines and carry over amounts, is set by the auctioneer. In almost all cases all bidders bid the maximum spread.

FIGURE 3. Initial Market Midpoint



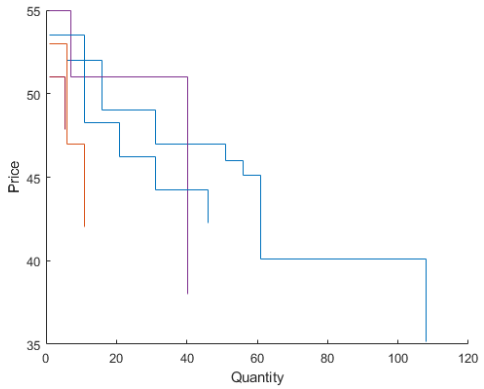
A buyer of CDS may have incentive to manipulate their price quote downwards in order to decrease the price floor/cap. However, the presence of fines and the fact that outlier bids do not get included in the average, should discourage this type of behavior. It has also been suggested that this may be used as signalling, similar to the LIBOR misquoting that has already been documented by Bonaldi (2017). Unlike the LIBOR context, where the individual quotes are revealed, only the average is revealed between rounds, limiting the signalling benefits. This substantially limits the ability to signal using this quote.

2.2.3. *Uniform Price Auction.* After the initial submissions, the initial market midpoint the size and direction of the open interest (quantity to be bought/sold in the auction) are announced. The market is then given between 30 minutes and two hours to incorporate this information. Next, a uniform price auction is held to clear the excess quantity. The bids submitted in the example auction are plotted in panel A of Figure 4. These bids are then summed to calculate a demand (or supply) curve and the point where it intersects the total quantity to be sold determines the clearing price as shown in panel B of Figure 4.

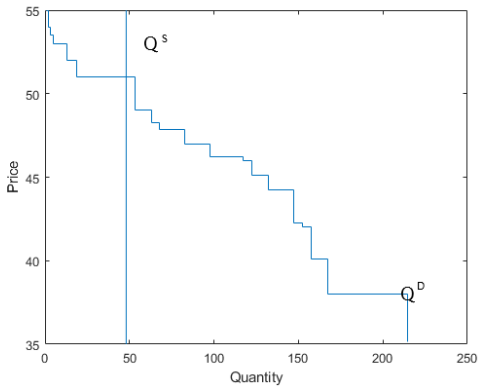
When bidding in this stage, the dealer chooses a demand curve to submit. Uncertainty about opponents' values leads participants to bid strategically. When deciding on the bid, each bidder considers the distribution of residual supply curves (representing the excess supply at each given price after accounting for the orders of opposing bidders). Knowledge of the NOI provides each bidder with information on the location of the aggregate supply curve and because the NOI results from the initial quotes it also informs them about the opponents' signals. Unlike a standard multi-unit auction, the final price is paid for all bonds acquired in the auction, for the initial quantity commitment and for all CDS contracts. The CDS positions provide dealers with an additional incentive to shade more or less by changing the effective number of units on which they pay the final price. In the current auction design, the first-round submissions allow some bidders to decrease their exposure to the auction price, which should help reduce the heterogeneity in exposure

FIGURE 4. Demand Function

(A) Bidder Demand



(B) Aggregate Demand



across bidders. Finally, the fact that the second round only allows bidding in one direction may constrain some participants from expressing their demand.

2.3. **Distinguishing Between Common and Independent Private Values.** Most empirical work on auctions requires the economist to make a modelling assumption on the information structure of the game, for tractability, given the complicated dynamic, multi-unit setting, this choice is limited to either the independent private values framework (IPV) or one based on common values (CV). In the CDS context there are factors which could lead both of these assumptions to be reasonable, and in theoretical work both have been used (Du and Zhu (2017) and Chernov et al. (2013), respectively). Specifically, IPV may be reasonable if the first-round price quotes effectively aggregate the common information held by different dealers and the remaining variation in values was driven by bidders' own costs of holding bonds, expectations of their own customer order flows, their value of liquidity, or their expertise in managing the complicated legal process of restructuring/liquidation, or their cost of holding bonds through the recovery process. On the other hand, common values is a reasonable assumption if there is a liquid resale market where these inventory/management costs are negligible.

In this section I empirically test for the presence of several correlations which are predicted if common values play an important role in the second stage strategic bidding decisions but which do not occur under IPV. I first perform the test proposed by Gupta and Sundaram (2015), which uses the variance of initial round quotes as a proxy for uncertainty. Due to the Winners Curse, when this uncertainty is high, bid shading in the second round should increase under common values.[28] I find no significant correlation between these measures (results can be seen in Table 5). I then focus on the 'independent' piece of the assumption and provide evidence in Figure 6 that the bids of two randomly

[28]Unlike Gupta and Sundaram (2015) there is no evidence of this relationship in my sample. I include controls for the first stage quantity, and the size of the NOI as these would introduce correlation even in a private value setting.

TABLE 5. Gupta and Sundara Test for the Winner's Curse

This table reports results from regressing the average slope of a bidder's stage two bid, against a proxy for the winners curse. The presence of the winner's curse suggests steeper stage 2 bids. The regression also controls for the bidders first stage quantity submission, the auction NOI, the N-steps submitted and a constant. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

	Slope of bids
IMM variance	-0.034
	(0.049)
Constant	0.154
	(0.097)

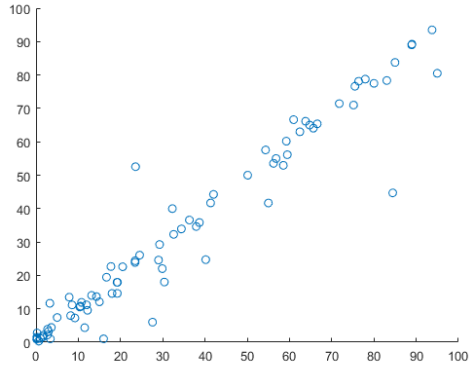
selected participants in the second stage of the auction are independent once the initial market quote is conditioned on. I formally test this relationship using a procedure proposed by Hickman et al. (2021). For each bidder I regress the own bid on the mean bid of opposing bidders. I find the average opposing bid does not predict each bidders own bid. Results are provided in Table 6. The test also suggests that unobserved auction heterogeneity does not play an important role. I then check if the auction price has any predictive power for post-auction prices after conditioning on the information available to bidders when submitting their round-two bids. Results of this exercise are presented in Table 7. I do not find any evidence of a correlation with bond prices 1, 5, or 30 days after the auction. Finally, I show that the bidders' own beliefs relative to the IMM level have no explanatory power for their second stage bid levels, suggesting that bidding in the two rounds does not reflect the same information. Results for this test are presented in Table 8.

Although there is likely to be some aspect of both private and common values in this setting, the results of these tests do not provide evidence of the Winner's Curse or of important within-auction correlation in bidder's values. This is consistent with bidders second-round bids, after they condition on the IMM quote being largely driven by idiosyncratic values.

FIGURE 6. Independence of Bids

Panels A and B present scatter plots of the raw and then the residuals of the quantity weighted average bid for 2 randomly selected participants in each auction. Under IPV these are uncorrelated. The correlation coefficient associated with this plot is 0.0139 and is not statistically significant. Auctions 43, 198 and 140 are dropped as they include outliers. These are auctions to sell the bonds, and these large bids are usually participants that sold into the auction putting a stop-gap bid in to ensure they will not receive less than that amount for the bonds they sold.

(A) Raw Bids



(B) Residualized bids

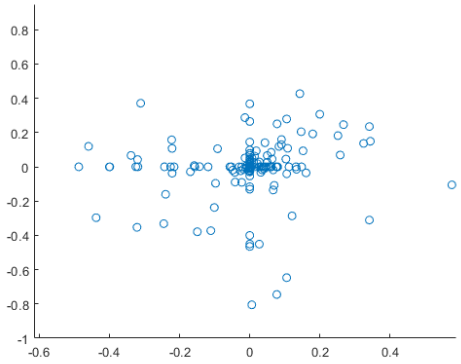


TABLE 6. IPV Regression Test

This table reports results from regressing the average bid for each bidder in stage 2 against the average bid by opposing bidders, controlling for factors that would explain across auction variation in the bid level in an IPV setting. A non-zero coefficient on the mean opposing bid would lead us to reject the null hypothesis of IPV. The specification follows the suggestion of Hickman et al. (2021) and adopts a cubic polynomial in N and the average opposing bids to control for variation in shading resulting from optimal bidding in an IPV model. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Controls	Own Bids
Mean opposing bid	0.047
	(0.9197)
NOI	-0.0132***
	(0.002)
IMM	0.812***
	(0.019)
Cubic polynomial in N_j	yes
Mean opposing bid x polynomial in N_j	yes
Constant	10.927
	(30.804)

TABLE 7. Post-Auction Prices

This table presents results from a regression predicting the post-auction price using the IMM price and the auction price. Results suggest that the final price is independent of the auction price. This is consistent with no information being revealed about the common value of the bond in that price. The price after each number of days is calculated as a volume weighted average and the sample is the set of auctions for which bond prices are available in TRACE. Prices are cleaned following Dick-Nielsen (2009). The securities missing price information include MBS, CMBS, and syndicated loans. Similar results are obtained when additional controls for bidding behavior are included. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Variable	Price after 30 Days	Price after 5 Days	Price after 1 Day
IMM price	1.739***	1.513***	0.525*
	(0.624)	(0.493)	(0.302)
Auction price	-0.771	-0.477	0.477
	(0.592)	(0.467)	(0.286)
constant	yes	yes	yes
N	56	56	56

TABLE 8. Second-Stage bids

This table reports results from regressing each bidders average stage 2 bid, against a measure of their initial beliefs about value (before announcement of the IMM). The regression also controls for the bidders first stage quantity submission, the auction NOI, the N-steps submitted and a constant. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Variable	Mean bid ($NOI < 0$)	Mean bid ($NOI > 0$)
IMM price	0.8356***	0.8586***
	(0.0128)	(0.0313)
$IMM_i - IMM$	0.9051	0.2898
	(0.5618)	(0.2645)
Controls	Yes	Yes
N	301	834

3. MODEL

3.1. **Players and Endowments.** The participants in the game are a set of dealers who are eligible to bid in the CDS auctions, $I_d \subseteq I$. This is the complete set of owners and sellers of CDS and owners of the underlying bonds. On the auction date, each dealer, i is endowed with a CDS position n_i. If $n_i \geq 0$, the dealer is a net buyer of protection while, if $n_i \leq 0$, the dealer is a net seller. Since these are derivative contracts, there is someone on each side of the position and $\sum_I n_i = 0$. Participants can also have any initial position in bonds, B_i. Note that for both the bonds and CDS positions these aggregation conditions hold over the entire set of market participants, not only the subset of dealers who bid in the auction.

Both the quantities B and n are denominated in hundreds of millions of dollars outstanding, so a bond payoff if no credit event occurs is $100B_i$ million. The final auction price is expressed as the cents on the dollar.

3.2. **Information.** Before making their choices, bidders receive independent draws of a vector $m_i = (s_i, n_i)$ from F_m, where s_i is a vector of private signals, and n_i is the one-dimensional positions in CDS contracts. The vector m is drawn from an absolutely continuous joint distribution with no holes and no mass points.[29] I assume bidders know the distribution F_m but not the individual draws of their opponents.[30] Let y_i denote the initial round quantity commitment of dealer i to purchase or sell bonds at the auction stage and let $v_i(q - y_i, s_i)$ denote the marginal value for the q^{th} unit of a bond purchased at auction.[31] I assume that these functions are bounded, weakly increasing in each component of s_i, and decreasing in q.[32]

In addition to the vector of private value-relevant private information, bidders receive a signal of the expected recovery value $\eta_i = R + \xi_i$, and $\xi_i \sim F_\xi$. These draws are IID across bidders and the bidders know the distribution F_η. I assume that this recovery value affects the bidders marginal value through a simple level shift: $R + v_i(q, s, B)$.[33] Bidders' initial belief of R is described by an uninformative prior. Both R and ξ are independent of the private signals and positions. Prior to the first-round bidding each bidder also receives orders from clients to submit physical settlement requests on their behalf. I assume that these orders arrive independently of all dealer's private information, and

[29]This assumption rules out that the position is a deterministic function of the vector of private information. While the bond or CDS positions of a dealer may reflect their private information, it seems reasonable that they are not perfectly correlated, for example: due to frictions in these markets leading up to the auction which may prevent some types of adjustment (eg. need to find the counterparty and agree on a price). The extent of OTC market frictions is well documented in theoretical work including Duffie et al. (2005) and Hugonnier et al. (2019) and in empirical work including Li and Schurhoff (2019) Bao et al. (2011) and Di Maggio et al. (2017). For a recent survey, see Bessembinder et al. (2020).

[30]This rules out that a bidders own position is informative of their opponents positions. This would be an important concern if, for example, bond holdings played a role in s_i and if one bidder owned most of the outstanding bonds and could therefore infer that their opponents held minor positions. This is not an important concern in this setting as the bonds owned by dealers usually make up a small share of the total outstanding, with large volumes of bonds owned by outside investors.

[31]This is indistinguishable from a model where bidder preferences depend on an initial position of bonds B_i which is also part of their private information (ie. $m = (s_i, n_i, B_i)$). In this case it would simply produce curves $v_i(B_i + q - y_i, s_i) = v_i(q - y_i, s_i, B_i)$. Without data on bond positions, the roles of signals and bonds in determining the marginal value cannot be separately identified, so I treat $v(q, \cdot)$ as the structural primitive and as a consequence simplify notation by writing $v_i(q - y_i, s_i)$ throughout.

[32]These bounds occur naturally in this setting, as no bidder should believe that the bond is worth a negative amount, and no bidder should believe that they will receive more returns than promised by the bond before the credit event.

[33]The choice of an additive R is motivated by the fact that the variance of bidders' information does not appear to shift with the level of R as a multiplicative form would predict. Evidence for this is presented in Section B.2.

each dealer receives a unique draw (y_i^c) from the distribution of these shocks Υ.[34] Finally, all bidders face a dealer-specific cost of submitting a bid in the first round, $c_{\kappa i} \sim \kappa$, and a (complexity) cost of submitting each step in the second stage of the game, $c_\iota \sim \iota$.

3.3. **Actions and Timing.** The bidders start with a quantity of bonds, a net position of CDS contracts, and some private signal indicating their private benefit from finishing the auction with q units. They also receive a signal about the expected recovery value η. At the start of the auction, each dealer receives orders to submit on behalf of their clients for physical settlement y_i^c.

Given their signals and position, dealers choose an initial round quantity (y_i) to commit to purchase/sell at the auction stage and a price quote that will be used to determine the IMM. In the first stage, bidders choose a quantity commitment y_i in the set:

$$\mathcal{Y}_i \equiv \{y_i | y_i \in [min(n_i, 0), max(0, min(B_i, n_i))]\}.$$

In keeping with the restrictions on participation from the auction rules, if $n_i \leq 0$, then $y_i \in [n_i, 0]$ while if $n_i \geq 0$, $y_i \in [0, n_i]$. The choice y_i is discrete with the interval of the minimum deliverable bond denomination. The total quantity submitted by each bidder is $y_i^o \equiv y_i + y_i^c$; the total submissions of other dealers is $NOI_i \equiv \sum_{j \neq i} y_j^o$; and the total submission of all dealers is $NOI \equiv \sum_{i \in \mathcal{I}_d} y_i^o$.

The total quantity orders from all the participants are announced along with the average price. After the first stage, bidders have information on the open interest, the initial market midpoint $\Omega = (NOI, IMM)$, as well as their own contribution to the NOI, y_i^o. Note that y_i^o influences the expected distribution of opponents' signals rather than y_i as this allows the bidder to deduce that the total submissions of opponents were $NOI - y_i^o$. Given this information, the bidder decides on a set of K_i pairs (b_{ik}, q_{ik}) to submit as a bid into the second stage uniform price auction.

In the second stage, bidders choose an action from the restricted set of strategies denoted by $\gamma(p|m_i, \Omega, y_i^o)$. This function describes the quantity γ allocated to bidder i at price p. The strategies γ_i for each player lie in the set of possible actions A_i, defined similarly to Kastl (2011):

$$A_i = \{(b_i, q_i, K_i) : dim(b_i) = dim(q_i) = K_i \in \{0, 1, 2, ..\bar{K}\}, b_{ik} \in \{0, 0.125, 0.25, ..., 100\},$$
$$q_{ik} \in [min(0, NOI), max(0, NOI)], b_{ik} \geq b_{ik+1}, q_{ik} \leq q_{ik+1}\}.$$

3.4. **Initial Market Price Quote.** I take the initial round price quote as a reflection of the value relevant information common to all bidders. I am therefore assuming that dealers do not use this quote to signal to other dealers their private values. This assumption

[34]It is important to note that learning y_i^c does not provide any information to the dealer about the bids an opposing dealers is likely to submit in stage 2. This earlier stage submission differs from learning from client orders that occurs when clients submit bids through dealers in the auction stage as in Hortaçsu and Kastl (2012)). This is because any information relevant to common or resale values is already conditioned out and since the customer order shock is independent, and dealer specific, it provides no information about the shocks of opponents. However, across-round learning about the level of competition still occurs in the model, through the announcement of the NOI.

simplifies the model and is reasonable in this setting because (i) only the average quote is reported before the second stage (rather than individual submissions), (ii) each bidder has limited ability to manipulate the averages and (iii) any attempt to engage in manipulation is likely to result in heavy fines and exclusion of the quote from the calculation of the average.

Assumption 1. *The p^{IMM} is the best guess of R for all the players after the first round results are announced and aggregates all the information in the common signals.*

Assumption 1 is key to guaranteeing that an equilibrium exists and allowing us to characterize the equilibrium behavior in the second-stage bidding game. It plays a helpful role in the empirical analysis by removing the common information components from valuations and provides an auction-specific measure of the bidders perceived values to help control for across-auction heterogeneity.[35]

Assumption 1 imposes two distinct simplifications. First, that bidders submit a truthful guess of the common component of values in the first stage, and second, that upon learning the p^{IMM}, bidders share a common expectation of the common value component.[36]

The trade-off is easy to understand: by decreasing their initial quote, a bidder will tend to decrease the IMM, which decreases the expected price floor or ceiling and may signal to opponents a lower expected resale value for the bond, leading to smaller opposing bids. However, by lowering their quote, dealers (i) increase the chance that their quote is not in the average (ii) reduce their bid that will be carried over into the second stage of the auction and (iii) increase the chance that they receive a fine. I argue that the costs from the fine discipline the distortions in the IMM.[37]

The Assumption 1 means that p^{IMM} is a sufficient statistic for the common part of bidders' values. This means that expectations in the second stage do not depend on the initial private signal, ξ, that inform bidders' first round quotes. Although this differs from the beliefs implied by Bayesian updating, this is likely to produce a similar set of beliefs. The updating should be similar because the second signal is much less noisy than the first (it aggregates the diffuse information from all participants), and the size of the correlation between the own submission and the outcome is difficult for bidders to evaluate. This simplified form of updating means that strategies in the second round no longer depend on the initial signal of the common component. This is consistent with the

[35]A central issue that arises when estimating demand systems is unobserved heterogeneity. We need to make sure that variation in quantity choices is attributable to variation in prices and not an omitted variable that is correlated with price, e.g. quality. These initial market quotes let us condition on the bidders' shared beliefs about the value and therefore capture differences from auction specific characteristics like quality.

[36]The requirement that the best guess of the common value is equal to p^{IMM} can be relaxed so that it is simply a monotone function of this price.

[37]Unlike the quoting game studied by Conley and Decarolis (2016) collusion in these quotes would be difficult to sustain because the direction that dealers want to manipulate the quote depend on their private CDS positions. If bidders were colluding we would expect to see a subset of quotes away from others—which we do not observe.

lack of correlation in the data between the first round quote and second round bid levels, as reported in Table 8.

3.5. **Stage 2: Auction Payouts.** In stage 2, bidders submit either a supply curve or a demand curve as appropriate to clear the open interest announced after the first stage. Because this submission occurs after learning the NOI, which is a function of opponents' private information, the players' expected distribution of opponents' signals in this stage will depend on the first stage strategies. In addition, the distribution of opponents signals that each player expects will differ due to their knowledge of their own contribution to the NOI, which, recall, is y_i^o. This distribution can be written as follows:

$$F_{m|\Omega,y_i^o} = \int_{[\underline{m},\bar{m}]\times[\underline{y}^c,\bar{y}^c]\times_{j\neq i}} 1(NOI - y_i^o = \sum_{j\neq i} y_j(\mathbf{m}_j,\eta,y_j^c) + y_j^c) \prod_{j\neq i} f(\mathbf{m}_j)\Upsilon(y_j^c)d\mathbf{m}dy^c.$$

Assume that these beliefs, $F_{m|\Omega,y_i^o}$, leave positive mass on every $m \in [\underline{m},\bar{m}]$, are absolutely continuous and have no holes and no mass points. I will show that these properties are satisfied such that beliefs are consistent with Bayesian updating given the equilibrium strategies. Therefore, these strategies and beliefs are a Perfect Bayesian equilibrium.

The auction choice is then to choose a strategy from the restricted set of strategies denoted by $\gamma(p|m_i,\Omega,y_i^o)$ which describe the quantity γ allocated to bidder i at price p. The strategies γ_i for each player lie in the set of possible actions A_i. The bidder chooses the strategy in γ_i in order to maximize the expected auction profits. Let the distribution of opponents' expected play given the information in Ω be denoted by L.

$$\Pi^A(m_i,y_i^o,L,\Omega) = max_{\gamma(.|m_i,,y_i^o,\Omega_i)} \int_m \int_0^q \Pi(m_i,b,q)dH(q,b|\Omega,m_i,\gamma(\cdot|m,y_i^o,\Omega))dL(m|y_i^o,\Omega) - \sum_{k=1}^{K_i(\gamma)} c_{ik}.$$

The bidder's profits in the auction is made up of three components: (i) the cash settlement on their existing CDS positions — paid at the auction clearing price, (ii) the auction payments — made for the quantity bought in the auction plus the commitment from the first round, and (iii) the benefit from the bonds bought/sold in the auction. I rewrite this problem as follows:

$$max_{\{b_k,q_k\}_{k=1}^{K_i},K_i} \sum_{k=1}^{K_i} \int_{b_{k+1}}^{b_k} [\underbrace{(100-p)(n_i)}_{\text{cash settlement}} + \underbrace{[R+v(q_k-y_i,s_i)]q_k}_{\text{Benefit from final bonds}} - p\underbrace{(q_k-y_i)}_{\text{quantity}}]f(p|NOI,p^{IMM},y_i^o)dp - c_{ik},$$

where c_{ik} is the cost of submitting each step and defined in section 3.2.

Optimality of the chosen bid implies the set of first order conditions (FOC) for demand bids in Equation 1.[38] Note that when a tie occurs the quantity is split pro-rata. This

[38]This FOC does not account for any bounds on the price. In reality the price in an auction to buy, is bounded on $[0, p^{IMM} + 2 * spread]$ and in an auction to sell on $[p^{IMM} - 2 * spread, 100]$. This may lead to

gives one equation for each step k with each derived by considering perturbations in the quantity q_k at the given price step. In any BNE, for almost every s_i, every step k in the K_i step function must satisfy the following equation.[39]

$$\begin{aligned}
&Pr(b_k > P^c > b_{k+1}|y_i^o, \Omega)[R + v(q_{ik} - y_i, s_i) - E_{M_{-i}|m_i}(P^c|b_{ik} > P^c > b_{ik+1}, y_i^o, \Omega)] \\
&Pr(b_k = p^c \wedge Tie)E[(R + v(q(S, \gamma(\cdot|S)) - y_i, s_i) - b_k)\frac{dQ^c}{dq_k}|P^c = b_k \wedge Tie] \\
&Pr(b_{k+1} = p^c \wedge Tie)E[(R + v(q(M, \gamma(\cdot|M)) - y_i, s_i) - b_{k+1})\frac{dQ^c}{dq_k}|P^c = b_k \wedge Tie] \\
&= (q_k + n_i - y_i)\frac{\partial E[p^c; b_k \geq p^c \geq b_{k+1}, y_i^o, \Omega]}{\partial q_k}.
\end{aligned} \tag{1}$$

Simplifying to remove ties and collecting $\alpha_i = y_i^o, \Omega_i$:

$$\begin{aligned}
&Pr(b_k > p > b_{k+1}|\alpha_i)[R + v(q_{ik}, s_i) - E_{m_{-i}|\alpha_i}(P|b_{ik} > p > b_{ik+1}, \alpha_i)] \\
&= (q_k + n_i - y_i)\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q_k}.
\end{aligned} \tag{2}$$

A similar argument can be applied for the case of bids to supply bonds, leading to the equation:

$$\begin{aligned}
&Pr(b_{ik-1} > p > b_{ik}|\alpha_i)[-R - v(q_{ik}, s_i) + E_{m_{-i}|\alpha_i}(P|b_{ik-1} > p > b_{ik}, \alpha_i)] \\
&= (q_k + n_i - y_i)\frac{\partial E[P; b_{ik-1} \geq p \geq b_{ik}|\alpha_i]}{\partial q_k}.
\end{aligned} \tag{3}$$

To simplify expressions in the following sections I focus on the case of excess supply. All the expressions are easily adapted to the case of excess demand.

Equation 1 is very similar to the FOC derived in Kastl (2011) and the FOC for an oligopolist with uncertain demand as in Klemperer and Meyer (1989), with the important additions of the price impact from cash settlement and initial quantity commitments. The LHS of the equation represents the marginal cost of quantity shading: the difference between the marginal utility and the expected price; while the RHS represents the marginal benefit of quantity shading from the savings on the inframarginal units. There are two important differences relative to Kastl (2011): (i) bidders learn about the expected level of competition and the total supply based on the NOI and so the expectations condition on this outcome, (ii) the CDS position, less quantity commitments, influences the importance of the price savings from quantity shading. The key difference is that the CDS position changes the number of units on which the bidder pays the market clearing price.

corner solutions where one dealer purchases all the quantity at the highest possible price or sells all the quantity at the lowest possible price, simply to influence their CDS payout. Whether this is a concern in practice depends on the support of $(n - y)$. The corner solution does not occur often in the observed bidding data and so I ignore this case in the discussion.

[39]This result is derived in Kastl (2011) for the standard multiunit auction. Once I apply the conditioning described above, my model is a special case of this game, where the expected payment term is transformed to be $E_{S_i}[(Q^c(S_{\neg i}, s_i) - n_i - y_i)P^c(S_{\neg i}, s_i)]$ units instead of q_{ik} in that setting.

For buyers of CDS it increases the number of units for which they must pay the price. This makes the bidder much more sensitive to any price changes that they may cause — leading them to shade their bids to buy more aggressively (or making them willing to supply more at lower prices).[40]

3.6. **First Stage Quantity.** The first stage quantity choice involves many strategic considerations. First, it changes the bidders exposure to the auction clearing price. Second, it changes the total quantity for sale in the auction (altering the distribution of marginal values to clear the market). Finally, it affects the expected level of competition for i's opponents due to its impact on the announced quantity.

The bidder chooses a quantity of bonds from $y_i \in \mathcal{Y}_i$ in order to maximize their expected profits from the auction: $max_{y_i \in \mathcal{Y}_i} E[\Pi^A(m_i, y_i + y_i^c, \Omega, L)|m_i, \eta_i]$.

Assumption 2. *y_i^c, the set of customer order shocks, is independent of the dealers own position n_i and has full support on the set of possible NOI.*[41]

Assumption 3. *Each dealer draws a cost κ of submitting a nonzero y_i. The support $Supp(\kappa)$ includes costs that satisfies the following. $\exists\ \delta_n > 0$ such that $\forall\ n_i \in [-\delta_n, \delta_n]$ there exists an open set of signals $\tilde{s}_i$, for which $\forall\ y_i \in \mathcal{Y}_i$, $\exists\ \Delta > 0$ such that $max_{\delta \in \{\delta || \delta| \leq |\Delta|\}} \Pi(y_i + \delta) - \Pi(y_i) \leq \kappa_i \in Supp(\kappa)$.*

This assumption is obviously satisfied if the cost distribution has an unbounded support. The weaker condition in the assumption is required to guarantee that for some positive mass of signals (with net CDS positions sufficiently close to zero), it is optimal for them to choose $y_i = 0$ for any customer order shock that they receive. This assumption guarantees that there exists a positive mass of dealers who pass through the customer order shocks they receive directly to the NOI. These shocks then smooth any possible jumps in the equilibrium distribution of NOI and insure that it has positive mass on its entire domain. The smooth distribution of NOI means that the maximum of these jumps is zero, and so the assumption is satisfied with any positive cost of submitting an initial quantity.

Proposition 1. *$\forall NOI \in [\underline{NOI}, \bar{NOI}]$, the probability density function $f_{NOI}(NOI) > 0$ and is continuous.*

Proof. By assumption, the cost of submitting is greater than the the largest jump in profits between two neighbouring choices of y. This implies that there is some positive mass (some interval in n near n=0) of signals whose optimal first round choice is $y_i = 0$. The shocks from client orders which are passed through directly to the initial quantity choice then mean that the density of NOI is continuous, with full support, which then implies that expected profits are continuous in y_i. □

[40]The existence of equilibrium is discussed in Appendix A. The existence of equilibrium in multiunit uniform price auctions with restricted strategy sets is an open question, however in the standard setting Kastl (2012) proves the existence of an epsilon equilibrium. That result does not apply to the CDS setting. I follow Kastl (2011) and impose a fine discrete grid of price levels. This is the case in practice as bidders can only express their prices to the nearest 1/8th of a cent.

[41]This assumption is slightly stronger than required to guarantee that an equilibrium exists, but some limit on the correlation is required and independence is required in the identification argument.

The presence of directly submitted customer order shocks means that from the perspective of a bidder it is never possible to distinguish which part of the observed open interest arose from a particular quantity submission by the bidders and which part from the pure random shocks. The addition of this pure additive noise means that it is impossible for the bidder to rule out any vector $\boldsymbol{m}$ from the observed NOI. Updating consistent with the continuous NOI distribution using Bayes rule then implies that the distribution of private information $F(s_0, n_0, s_1, n_1, ..., s_N, n_N|\Omega, y_i^o)$ is an atom-less distribution with common support and density f. This then satisfies the assumptions made on these beliefs in Section 3.5, therefore an equilibrium exists with these beliefs. Given this, the first stage is an incomplete information game, with continuous payoffs and so there exists an equilibrium.

The initial round quantity submission provides a method for making commitments that adjust the bidder's position and desire to strategically bid in the auction by shifting their exposure to the final auction price. These features are related to those of sequential markets in Treasury, or electricity market settings as in Allaz and Vila (1993), and Ito and Reguant (2016). Unlike in these forward markets, the initial round in this game is settled at the final auction price rather than a separate forward market price. Because of this, if the second stage game allowed for both supply and demand bids following the initial round, the initial round quotes would be non-binding and non-informative. However, the single directional second round means that sometimes a bidder will be unable to adjust the change in position from their initial quote and this provides a cost for a particular choice in the first round. I examine the impact of this alternative commitment cost on market power empirically using the model estimates.

4. Identification

I want to identify the joint distribution of marginal value curves and CDS positions, the distribution of entry costs for each additional step and the distribution of customer order shocks. I argue that all these distributions are set identified.[42] Although additional restrictions on the shape of $v()$ can greatly simplify the identification discussion, in this section I provide intuition for how the data restrict the sets of possible distributions without the use of functional-form restrictions.

The main identification argument uses a GPV-type approach (Guerre et al. (2000)) to estimate the bidders' marginal values for additional units that rationalize each observed bid. Unlike in GPV, or the standard multiunit auction case, where the unobservable value can be written as a function of observables, the unobservable value in a credit event auction depends on the level of the CDS position. That is, for any CDS position

[42]I do not separately identify the signals and bond position in the function $v()$. Doing this would require additional structure on this function. As there is a secondary market for bonds, there are not meaningful constraints from the bond position: ie. a dealer could sell more bonds than they owned by going to the market and buying more. This would enter the model as a shift in the value of selling which I estimate and so the estimated $v()$ should provide sufficient information to understand both factual and counterfactual bidding.

there is a unique unobserved value that rationalizes the observed bids. Imposing that marginal values are monotone decreasing eliminates all the CDS positions which imply nonmonotonic marginal value curves, leaving us with a joint set of CDS positions and marginal value functions which may be consistent with the behavior of each bidder.

4.1. **Marginal Value and CDS positions.** To begin, I show that a curve

$$\tilde{v}(q) = v(q, s_i, B_i) + (n - y)\frac{\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q}}{Pr(b_k > p > b_{k+1}|\alpha_i)}$$

is identified at the subset of quantities where steps are submitted. As in Kastl (2011), the terms $Pr(b_k > p > b_{k+1}|\alpha_i), E_{m_{-i}|\alpha_i}(P|b_{ik} > p > b_{ik+1}, \alpha_i)]$, and $\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q_k}$ are directly identified from observed bidding data. The common value term, R, is identified from Assumption 1 and initial quote submission data. Rearranging equation 3, gives the newly defined curve $\tilde{v}(q)$ as a function of identified objects.

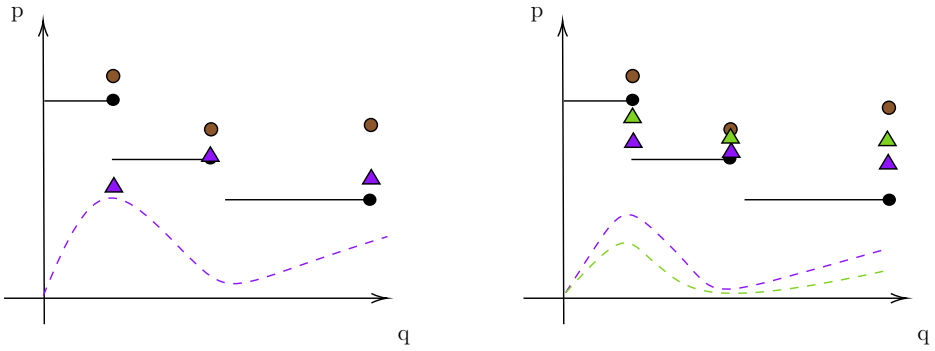
$$\begin{aligned}\tilde{v}(q) &= v(q, s) - (n_i - y_i)[\frac{\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\Omega, y_i^o]}{\partial q_k}}{Pr(b_k > p > b_{k+1}|\Omega, y_i^o)}] \\ &= E_{m_{-i}|\Omega, y_i^o}(P|b_{ik} > p > b_{ik+1}, \Omega, y_i^o) - R + (q_k)[\frac{\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\Omega, y_i^o]}{\partial q_k}}{Pr(b_k > p > b_{k+1}|\Omega, y_i^o)}].\end{aligned}$$

Given knowledge of the curve $\tilde{v}(q)$ as well as the ratio of the price impact $\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q}$ to the probability of clearing $Pr(b_k > p > b_{k+1}|\Omega, y_i^o)$, and the monotonicity and boundedness (assumed in the structure of the model) of the $v(q, s_i)$ allows us to bound the $v(q)$ and the possible $n - y$ simultaneously. That is: for $q_k > q_{k-1}$ it must be that $v(q_{k-1}) \geq v(q_k)$. If $\frac{\partial EP}{\partial q}$ is not monotone across the set of q_k where the curve $\tilde{v}(q)$ is observed, this provides an upper and lower bound. Intuitively, $n - y$ must be such that the observed changes in $\tilde{v}$ can be rationalized with $\frac{\partial EP}{\partial q}$ and a bounded, monotone decreasing function.

As an example, Figure 8 draws in black an observed bid curve defined by the set of steps. The dashed lines denote the observed ratio of price impact to win probability multiplied by different factors $(n - y)$. The round dots above the observed bids denote the $\tilde{v}(q)$, calculated from the observed price impact, probability and expected clearing price. The triangular dots show the marginal value curve associated with a particular level of $(n - y)$. That is, the triangles are defined so that the sum of the triangle and dashed line give the round dots $(\tilde{v}(q))$. In the left panel the implied marginal value curve is not monotone decreasing. This allows us to conclude that the $(n - y)$ factor is too large and cannot be part of the identified set. The right panel illustrates two possible $(n - y)$. The dashed curves (in indigo and green) illustrate the ratio of price impact to win probability multiplied by each of these factors respectively. The green triangles show the implied marginal value curve associated with the green dashed curve and the indigo triangles the marginal value curve implied by the indigo dashed curve. Because both the green and indigo triangles are monotone decreasing, the associated $(n - y)$ are part of the identified set.

FIGURE 8. Bounds from Monotonicity

The black lines denote observed submitted bids. The dashed lines denote the observed ratio of price impact to win probability multiplied by different factors (n-y). The round dots denote $\tilde{v}(q)$. The triangular dots show the implied marginal value curve: the sum of the triangle and dashed line give the round dots ($\tilde{v}(q)$).In the left panel the $(n-y)$ factor is too large and the implied marginal value curve is not monotone decreasing. The right panel illustrates two possible marginal (n-y) factors. The dashed curves (in indigo and green) illustrate the ratio of price impact to win probability multiplied by each of these factors. The green triangles show the implied marginal value curve associated with the green dashed curve and the indigo triangles the marginal value curve implied by the indigo dashed curve.



The bounds on $\tilde{v}$ also help restrict the set of $(n-y)$ that are consistent with the observed bids. For example if $\tilde{v}(q_s) \leq 0$, the fact that $v(q_s) \geq 0$ implies $(n-y)\frac{\partial E}{\partial q} \leq \tilde{v}(q_s)$ and for the upper bound, that $100 + (n-y)\frac{\partial E}{\partial q} \geq \tilde{v}(q_s)$. This set of restrictions is quite informative, as $n-y$ is constant across all quantity levels and many bids contain more than one step, which share the same $n-y$.

The information content of this identification argument depends on the observed differences in the value curve and price impact of shading across quantity levels, which all share the same $(n-y)$ within a given bidder. Take two quantity levels $q_1 < q_2$ at which bidder i submitted bids.

$$\tilde{v}(q_1) - \tilde{v}(q_2) = v(q_1, s_i) - v(q_2, s_i) + (n-y)\left(\frac{\partial E[P; b_1 \geq p \geq b_{1+1} | \alpha_i]}{\partial q_1} - \frac{\partial E[P; b_2 \geq p \geq b_{2+1} | \alpha_i]}{\partial q_2}\right).$$

The LHS of this equation is observed, as is the term inside the final set of brackets. By monotonicity of the marginal value curve, the difference $v(q_1, s_i) - v(q_2, s_i) \geq 0$ is known. If the difference in $\left(\frac{\partial E[P; b_1 \geq p \geq b_{1+1} | \alpha_i]}{\partial q_1} - \frac{\partial E[P; b_2 \geq p \geq b_{2+1} | \alpha_i]}{\partial q_2}\right) \geq 0$ then this provides an upper bound for $n-y$ while if it is negative then it provides a lower bound. Lets consider two pairs of points, with one pair providing an upper and the other pair the lower bound. The true value is given in the following expression.

$$(n-y) = \frac{\tilde{v}(q_1) - \tilde{v}(q_2) - (v(q_1, s_i) - v(q_2, s_i))}{\left(\frac{\partial E[P; b_1 \geq p \geq b_{1+1} | \alpha_i]}{\partial q_1} - \frac{\partial E[P; b_2 \geq p \geq b_{2+1} | \alpha_i]}{\partial q_2}\right)}$$

which can be decomposed into the observed (first term: difference in $\tilde{v}$ and the unobserved but bounded second term: difference in $v(q)$).

4.2. **Entry costs and Client orders.** I construct bounds on the distribution of client orders (y^c) by leveraging the constraints that bidders with net long (short) CDS positions remaining in the auction who submit physical settlement requests to buy (sell). Since y^c are independent of the original position, this allows me to identify the distribution of clients order shocks. To begin, take the bidders with $(n_i - y_i) > 0$ (and therefore $y_i \geq 0$). For these dealers, $y^c \leq y^o$. This means that the distribution of y^o on this subset of bidders gives a valid upper bound on the client order shock distribution: $Pr(y^o \leq y | (n_i - y_i) > 0) \geq Pr(y^c < y | (n_i - y_i) > 0)$. Similarly, for dealers with $(n_i - y_i) < 0$ (and therefore $y_i \leq 0$), it must be the case that $y^c \geq y^o$. This means that the distribution of y^o on this subset of bidders is a lower bound for the distribution of customer shocks. Because the y^c shocks are independent of n_i, the only selection in calculating the unconditional distribution comes from the mass of bidders where the sign of n_i cannot be inferred. This occurs if (i) the bidder chooses to submit their entire position $n_i = y_i$ (ii) $0 \in [\underline{(n_i - y_i)}, \overline{(n_i - y_i)}]$. Fortunately this mass is observed and so by adding it to the upper bound from the selected sample we obtain an upper bound on the distribution of client shocks. Finally, the bounds on $n_i + y_i^c = (n_i - y_i) + y_i^o$ together with the distribution of y^c provide bounds on the distribution of n, conditional on each curve $v(q, s_i)$.

The distribution of costs for submitting an additional step ι can be bounded from above by calculating the maximum profit difference a bidder could achieve by adding an additional step, and from below by comparing the true profit to expected profit with one less step. I do not consider identification of κ as it plays no role in the counterfactuals.

The discussion so far showed that the model primatives are identified conditional on choosing a non-zero number of steps. However, this leaves a problem of selection on observables. This, however, can be easily corrected. First note that for every signal draw there is a positive probability of submitting at least one step, as variation in the auction reverses the set of bidders most likely to be excluded. Further, at each signal vector m_i, I can calculate the probability of submitting $K_i = 1$ steps instead of zero by using the expected differences in profits from including the step. Since the distribution of costs is already identified, we can compare these profit differences to the costs to calculate the probability of submitting zero steps at each m_i.

5. Estimation

Despite being non-parameterically set identified, a fully nonparametric estimation would require far more data than are currently available. Therefore, I impose some parametric restrictions to reduce the dimensionality of the problem. I perform tests supporting many of these assumptions in Section 6.4. There are several important challenges for estimation of this model: (i) the model is dynamic, (ii) dealers have both private information and private positions, and (iii) it is common to submit only a smaller number of steps.

To begin, assume that the marginal valuation curves of each bidder are linear. That is, the marginal value can be represented by (i) a signal reflecting the value for the first unit of bonds acquired in the auction (ii) a rate at which the marginal benefit from each additional unit declines.

Assumption 4. *The dimension of the private signal is 2 and the form of the marginal value is linear* $v(q, s) = s_1 - s_2 q$.

This implies that for all bidders who place more than three steps, the linear restriction is over-identified and therefore these cases can be used for testing. In Section B.6 I show that the R-square from the linear fit is high and that the addition of a quadratic term does not result in a large change in either the R-square or model estimates.

Assumption 5. p^{IMM} *is a sufficient statistic for the observed (across auction) variation in bond traits* Z *and these traits only impact* R*, not the joint distribution of* s_1, s_2, n.

When estimating the distribution of opposing bids that a bidder expects to face, we need to condition on Ω, y_i^o and the observed characteristics of the set of bonds eligible for submission to the auction. The observed initial market quote picks up a large amount of the across auction heterogeneity, including differences due to the observable bond traits and those that are observable to bidders but not the econometrician. However, the bond traits and volumes may affect values beyond their role in determining the IMM. I test for this and show that they have no explanatory power for auction outcomes once the IMM price is conditioned on. The supporting regressions are presented in Section B.5. Given this evidence, I assume that the IMM is a sufficient statistic for capturing observable differences on bond traits, Z. This assumption reduces the dimension of the estimation problem, which improves the power of the estimates.

I parameterize the distribution of s_1, s_2 and n using 4-, 4- and 6-parameter cubic B-splines, respectively, to describe the quantile functions of the marginal distributions and impose that the correlation structure is given by a Gaussian Copula.[43] I parameterize the distribution of entry costs as Normal and estimate the mean and standard deviation. Finally, I specify the customer order shocks distribution as Normal, with mean zero and estimate the variance.

In the previous section I showed that the model is non-parametrically set identified. I now provide intuition for why the model that I estimate with these additional restrictions is point identified. First, for every bidder that submits three or more steps we learn a unique $(n - y), s_1, s_2$. For any combination of $(n - y), s_1, s_2$ we also know the difference in profits from using $K = 1, 2, 3, ...$ steps. By comparing the probability of submitting K steps when the difference in expected profits are some fixed level, we can identify the probability of a submission cost exceeding/not exceeding that level. Since the submission cost distribution goes from $[0, \infty)$ and the change in expected profits are weakly positive, then for any draw of $(n - y), s_1, s_2$ the bidder will sometimes submit three or more steps and so the probability of that vector is known. For each bidder we also observe the y^o. As in the previous section we can construct bounds on this distribution using restrictions on the eligible submissions of the bidder. This does not guarantee a unique σ_{y^c}. However, if we assume that the distribution of y_i has a compact support then together with the

[43]The Gaussian Copula facilitates the quick generation of correlated random draws for simulated integration during the estimation.

normality of the errors y^c results from Bertrand et al. (2019) insure identification given y^o.

The estimation contains three distinct steps. In the first, I use techniques developed in the literature on multi-unit auctions, and use a weighted resampling estimator to estimate $Pr(b_k > p > b_{k+1}|\Omega, y^o)$, $E_{m_{-i}|\Omega,y^o}(P|b_{ik} > p > b_{ik+1}, \Omega, y^o)$ and $\frac{\partial E[P;b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q_k}$, where weights are used to control for selection on observables as well as other behavioral responses of bidders to these observables. In the second step I estimate functions which approximate the differences in profits for a given bidder of bidding using 0,1,2, or 3 steps. In this way I can control for selection (arising from the fact that the number of steps reveals additional information about the values).[44] Conceptually, this calculation could be made inside of the final step, however nesting this calculation is not computationally feasible. In the final step, I combine the estimates of these components, with the restrictions from the first order conditions, to form a set of moment conditions which allow for the parameters of the joint distribution of s_1, s_2 and n and the parameters of the entry costs and customer order shock distribution to be jointly estimated. The next sections discuss each of these components in detail.

5.1. **Stage 1: Resampling.** All the terms in the bidder's FOC are functions of the three terms: (i) $Pr(b_k > p > b_{k+1}|\Omega, y^o)$ — the probability of being allocated quantity q_k associated with price bid b_k, (ii) $E_{m_{-i}|\Omega,y^o}(P|b_{ik} > p > b_{ik+1}, \Omega, y^o)$ — the expected clearing price conditional on winning q_k and (iii) $\frac{\partial E[P;b_k \geq p \geq b_{k+1}|\alpha_i]}{\partial q_k}$ — the price impact of increasing q_k.

In the first stage I therefore construct estimates of

$$\begin{aligned} &E_{M_{-i}|m_i}(P^c|b_{ik} > P^c > b_{ik+1}, y_i^o, \Omega) \\ &\frac{\partial E[P^c; b_{ik} > P^c > b_{ik+1}, y_i^o, \Omega]}{\partial q_{ki}} \\ &Pr(b_k > P^c > b_{k+1}|\Omega, y_i^o). \end{aligned} \tag{4}$$

This estimation step follows directly from Hortaçsu and McAdams (2010) and Kastl (2011). To handle shifts in bids due to observable differences across auctions, Hortaçsu and McAdams (2010) propose a conditioning approach weighting by the traits in the resampling process used to approximate Equation 4. With this approach, weights are used to control for both selection and behavioral responses to observables Ω, y^o. I allow the conditioning on NOI and the own physical settlement request submitted by resampling from the opponents of bidders with similar NOI of others and own requests to enter non-parametrically in the estimation by using kerenel weighting of these traits across auctions.

The challenge in this setting is that Ω includes p^{IMM}, NOI, y_i, so the kernel weights must reflect the similarity of the information set faced by individual bidders. (That is, I

[44]For example, when a bidder submits two steps the FOC provides a set of possible m_i, but some points in that region are very unlikely to bid only 2 steps and others and this must be accounted for in the aggregation.

need to construct an estimate of $Pr(Bidder_j|\Omega_i)$.) To do this I use the logic that bidders with information sets that are similar, should expect similar opposing bids.

The resampling scheme should put the most weight on an opponent showing up that looks like the opponents of a bidder with a particular information set. For example, if bidder 1 in auction 1 and bidder 3 in auction 15 have the same information sets, they should expect to face opposing bids from the same distribution of opponents' bids. To evaluate this in a tractable way, begin by finding the bidder with the most similar information set to bidder i, in each other auction. For each of these most similar bidders measure the difference between their information sets, and, using this distance, define an auction level weight that will be applied to all the opponents of that most similar bidder, while give zero weight to resampling the single most similar bidder. In this way two bidders with the same information set should expect to face the same set of opponents.[45] This gives a set of weights:

$$w_{Aj} = \begin{cases} (\sum \frac{\max_{l\in A} K(\frac{\alpha_l - \alpha_i}{bw})}{\sum \max_{l\in A} K(\frac{\alpha_l - \alpha_i}{bw})})/\mathcal{I}_{dj} & l^* \neq j \\ 0 & l^* = j. \end{cases} \tag{5}$$

Asymptotically this is consistent because as the size of the bandwidth shrinks, only opposing bidders from auctions where the most similar bidder had the same information set receive positive weight. Note that nothing in the information set is estimated; these components are all observed without error. Implementing this in practice I resample from the quantity and price shares, which helps avoid extreme draws. This normalization has no effect asymptotically, because as the bandwidth shrinks samples are drawn from auctions with identical p^{IMM}, NOI.

5.2. **Stage 1b: Selection.** In estimating the second stage of the model, it is important to incorporate the bidders who submit less than three steps, despite the fact that the signals and private position $[s_1, s_2, n]$ that rationalizes their observed bid cannot be uniquely pinned down. For bidders that use less than three steps, there are three unknown values to estimate but less than three observed points. Rather than a unique vector of private information, therefore, the restrictions from the FOCs give us a set of signals and positions that could be consistent with the observed bid.

Each bidder decides how many steps to use in their bid function by comparing the additional expected profits from including another step to the cost of submitting a bid with one more step. Because the differences in expected profits depend on bidders' private information, some $[s_1, s_2, n]$ are more likely to result in submissions with three or more steps. This means that the distribution of values conditional on having submitted three or more steps will be different from the population level distribution. To account for

[45]Eg. for bidder 1, in auction 1: the most similar may be bidder 4 in auction 2 and bidder 7 in auction 3 ect. I then weight the opponents of e.g. bidder 4 in auction 2, based on how similar bidder 4 is to bidder 1. So maybe bidder 4 in auction 2 is very similar to bidder 1 in auction 1 but bidder 7 in auction 3 is quite different. In this case the opponents of bidder 4 (ie. bidders 1-3,5-11) in auction 2, would be resampled with high probability while the bidders 1-6,8-11 in auction 3 would be resampled only very rarely.

this difference, we can compute the probability that any $[s_1, s_2, n]$ submits three steps by comparing the benefit of adding an additional step (which can be calculated as a function of observables) to the cost distribution for the individual-specific random cost of submitting an additional step.

To incorporate the bidders who submit less than three steps, I integrate over the set of possible values consistent with the observed bid while re-weighting to account for the probability that the bidder chose to submit less than three steps, which varies across the points in the set. To handle this I include in the integral an estimate of the probability of submitting $K = 0, 1, 2, 3+$ steps conditional on the level of $(s_1, s_2, n - y)$. For each bidder, it is difficult to compute the expected differences in profit from submitting their chosen number of steps rather than some alternative. To maintain tractability, I specify the differences in profit levels using a restricted functional form:

$$\Pi(3, (s_1, s_2, n-y), \Omega, Z) - \Pi(2, (s_1, s_2, n-y), \Omega, Z) = h_3((s_1, s_2, n-y), Z, \Omega, \beta) + u,$$
$$\Pi(2, (s_1, s_2, n-y), \Omega, Z) - \Pi(1, (s_1, s_2, n-y), \Omega, Z) = h_2((s_1, s_2, n-y), Z, \Omega, \beta) + u.$$

For the functional form of h_k I use a second order complete polynomial in $n, s_1, s_2, 1(NOI > 0), IMM$. The second order complete polynomial should allow for most of the important interactions between these variables (see discussion in Judd (1998)).

I then compute estimates of these equations by calculating the optimal bids with 1,2,and 3 steps for 1000 random draws of possible signal vectors, uniformly sampled between the bounds of the signals $[\underline{s}_1, \bar{s}_1]x[\underline{s}_2, \bar{s}_2]x[\underline{n}, \bar{n}]$ where the bounds are estimated using the set of bidders who submitted more than three steps in the original data (and hence for whom the signal vector is perfectly known). I assign each signal vector to an auction where the auction traits are chosen to be those from a randomly selected auction. Then, I compute these profit differences on that sample. I truncate the change in expected profits calculated (the dependent variable) at \$500M.[46]

5.3. **Stage 2: Aggregation.** In the first stage I obtained consistent estimates of coefficients in a linear system that bidders bids must satisfy. Depending on the number of steps submitted this system might be over, exactly, or under-identified. I then solve the optimal set of parameters using the simulated method of moments and use simulation to solve the difficult integrations (eg. the integration over the multiple solutions that satisfy the system of equations).

Before discussing the moment conditions, I revisit the linear system that is formed within a bidder by their set of K_i optimality conditions. This gives a system of equations,

[46]Truncating allows us to achieve a good fit at levels where the cost shock plays an important role (small profit differences). The fitted model still implies probabilities of submitting an extra step very close to 1 for truncated bidders.

where each step satisfies:

$$s_1 - s_2 q - (n_i - y_i)[\frac{\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\Omega, y_i^o]}{\partial q_k}}{Pr(b_k > p > b_{k+1}|\Omega, y_i^o)}]$$
$$= E_{m_{-i}|\Omega, y_i^o}(P|b_{ik} > p > b_{ik+1}, \Omega, y_i^o) - R + (q_k)[\frac{\frac{\partial E[P; b_k \geq p \geq b_{k+1}|\Omega, y_i^o]}{\partial q_k}}{Pr(b_k > p > b_{k+1}|\Omega, y_i^o)}].$$

For each bidder I calculate an estimate of the private positions $[s_1; s_2; (n-y)]_i$. By collecting the terms above and multiplying $[s_1; s_2; (n-y)]_i$, and rewriting this in matrix form gives

$$\hat{A}_i[s_1; s_2; (n-y)]_i = \hat{d}_i.$$

For all bidders the objects A, d are measured with error. For bidders with fewer than three steps, I simulate $(n-y)$ and so all the finite sample errors occur in the dependent variable. However, when three or more steps are submitted, the term which multiplies (n-y) is a regressor with measurement error. I adopt a minimum distance shrinkage approach to correct for these errors.[47]

I then solve the following set of moment conditions simultaneously. Standard errors are calculated using the bootstrap, where resampling is done at the auction level and a particular bootstrap draw is held fixed throughout the first stage resampling estimator, the selection estimation and the second stage.[48]

For each of the marginal distributions for levels of α at each decile, where m_j denotes the jth element of m_i, $x_{j\theta}(\alpha)$ is the inverse of the marginal distribution $F(x_{jn}|\theta) = \alpha$, and $\mathcal{M}_\rangle$ is the set of m, y_i for which $\hat{A}m = \hat{d}$ giving

$$E[\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} 1(m_j \leq x_{j\theta}(\alpha))1((m, y^0 - y^c) \in \mathcal{M}_i)Pr(K_i|\Delta\Pi(m_i, y_i,), \theta)h(y^c; \theta)f_\theta(m; \theta)dmdy^c - \alpha] = 0, \quad (6)$$

and a moment condition for the covariance

$$E[\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} (m-\mu_m)(m-\mu_m)'1((m, y^0 - y^c) \in \mathcal{M}_i)Pr(K_i|\Delta\Pi(m_i, y_i), \theta)h(y^c; \theta)f_\theta(m; \theta)dmdy^c - \theta_\rho] = 0. \quad (7)$$

To pin down the distribution of y^c I leverage the restrictions on y_i that each bidder can submit. These restrictions together with the observed y^o imply a set of possible submissions $\mathcal{Y}_y$. When combined with a y^o, each $y_i \in \mathcal{Y}_y$ is associated with some y_i^c, and it must be the case that when these sets are aggregated across bidders the implied

[47]In this application the problem is further complicated relative to the Empirical Bayes case, because the errors in this term are correlated with the measurement error in the dependent variable where A_{*3} is multiplied by q. The resulting bias is given by $\hat{[s_1; s_2; n-y]} = [s_1; s_2; n-y] + (A^TA)^{-1}A^T\epsilon_2(-(n-y)-q)$. To evaluate this bias, I calculate measurement error (ϵ_2) by bootstrap resampling of the first stage, and apply a correction by solving that equation. Note, that asymptotically this ϵ_2 vanishes and so even without the correction the estimates are consistent.

[48]The bias-correction factor for the bidders who submit 3 or more steps is held fixed across replications. It is estimated using 1000 bootstrap replications of the first stage, and it would be computationally infeasible to correct this on each sample. Further, the additional uncertainty from this term is likely to play only a very small role.

probability of being below some point $\tilde{y}$ lines up with the probability in the proposed y^c distribution.

$$E[\int 1(y^c \leq x_{y\theta}(\alpha))1(y^o - y^c \in \mathcal{Y}_y)h(y^c;\theta)dy^c - h(y^c;\theta)] = 0$$

I also leverage the restriction that y_i has a compact support which I assume is given by the minimum and maximum observed holdings in the dataset of **?** and verify that the estimated minimum and maximum y_i are inside this support.

Finally, to pin down the parameters of the c_i distribution, I use the observed probability of submitting K steps, along with the observed differences in the profit functions to construct moments:

$$E[\Phi(\Delta\Pi_{32}, \hat{\mu}, \hat{\sigma}) - 1(K_i = 3)] = 0, \tag{8}$$

$$E[(1 - \Phi(\Delta\Pi_{32}, \hat{\mu}, \hat{\sigma}))\Phi(\Delta\Pi 21, \hat{\mu}, \hat{\sigma}) - 1(K_i = 2)] = 0, \tag{9}$$

$$E[(1 - \Phi(\Delta\Pi_{21}, \hat{\mu}, \hat{\sigma}))\Phi(\Delta\Pi 10, \hat{\mu}, \hat{\sigma}) - 1(K_i = 1)] = 0, \tag{10}$$

$$E[(1 - \Phi(\Delta\Pi_{10}, \hat{\mu}, \hat{\sigma})) - 1(K_i = 0)] = 0, \tag{11}$$

for three, two, one and the zero steps respectively, where $\Delta\Pi_{jk}$ denotes $\Pi(j, m_i) - \Pi(k, m_i)$, integrated over the possible vectors (m_i, y_i) with parameters θ as in the previous conditions (eg. Equation 6).

The model is basically a random effects model, with selection and censoring, and where the explanatory variables contain some measurement error. The simulation is performed over the integrals which are replaced by sums over S simulated draws.[49] Because this is a multi-dimensional joint distribution, I use importance sampling from the marginal distributions when integrating.

6. Results

6.1. **SMM Estimates.** The estimated parameters are presented in Table 9. The distribution of CDS positions is presented in tens of millions of dollars. The signal distribution is in terms of cents over or under the common value component. Note that the signal distribution is truncated in a way that is specific to the individual auction and that all three of the signal intercept, signal slope and bond position, as well as the initial quantity submission, interact to determine the actual effective intercept of the value curve (for the value of acquiring an additional unit in the auction). These actual effective intercepts are plotted in Figure 9 for three different values of the common value, R: 9, 33 and 80.

The distribution of CDS positions (n) implied by the estimation is fairly close to the distribution reported in Paulos et al. (2019) based on regulatory data on the positions of roughly half the dealers (those regulated by the Federal Reserve) in a sample of 15 of the CDS auctions from 2013-2017. The mean estimated position implied by the estimates is \$9.55M compared to \$8.23M in their regulatory data.

The estimated correlation between s_1 and n is negative. This is consistent with the incentive for bidders to hold too many CDS in order to avoid being constrained during

[49] In practice, I set S=1000 for the main estimation. Expanding to S=4000 instead should reduce the role of the simulation error by half. The resulting estimates are similar.

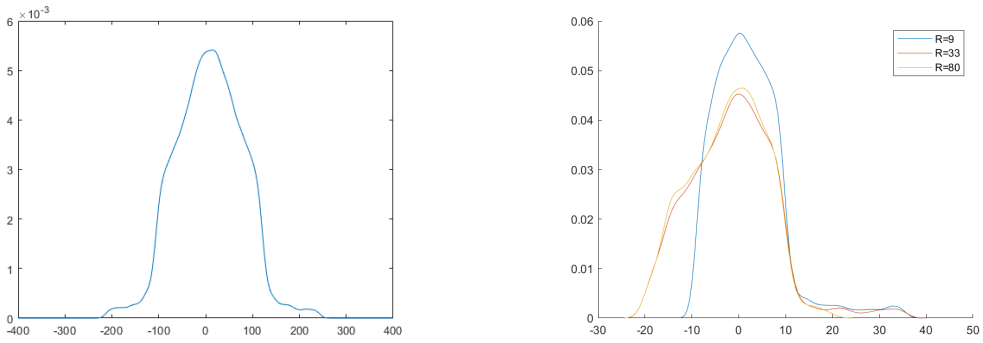
TABLE 9. Estimated Parameters

This table presents the coefficients for the spline quantile functions for the three marginal distributions and correlations are presented in this table. These are the results of estimating equations 6 and 7. Standard errors in parentheses.

Intercept	Slope	CDS position		Other Parameters
-16.5937	0.0005	-26.7647	Entry cost mean	-4.970
(3.2334)	(0.5913)	(4.7445)		(23.5679)
-0.7233	0.0005	-11.9472	Entry cost Std	16.5899
(1.4821)	(3.2783)	(2.7226)		(14.3125)
-0.7233	13.2292	0.7185	Client Shock Std	14.8309
(1.4916)	(3.9444)	(0.444)		(3.447)
13.3827		0.7185	Correlation: s1 s2	0.686
(3.6308)		(0.5251)		(0.0527)
36.595		13.7827	Correlation: s1 n	-0.67
(6.5187)		(3.7658)		(0.237)
		23.9233	Correlation: s2 n	-0.272
		(4.5737)		(0.216)

FIGURE 9. Marginal Distribution: CDS Positions

The left panel plots the estimated distribution of CDS positions (n_i). The right panel plots the estimated distribution of effective intercept shown at three levels of R, the common value component. The plots show kernel smoothed densities from 10000 simulated draws from the distributions implied by the quantile functions with parameters in Table 9.



the credit event auction process, as discussed in Du and Zhu (2017). It is also consistent with bidders with low post-default bond values buying more insurance in CDS markets.

6.2. **Expected Surplus.** In order to give some context to the estimates I compare the expected surplus and expected change in price that would result if bidders used truthful bidding, i.e. if bidders reported directly their implied value functions. This comparison removes the incentive to bias the price from the CDS contract position as well as from the competitive effects from information rents. The results integrate over possible draws of

the individual private information using 1000 simulated draws of potential bidders. The results of this calculation show that, on average, the prices are lowered by a median of 1.58 cents on the dollar, or mean of 4.00 cents on the dollar, as a result of market power in the auction. These results are similar to the gaps between the auction price and secondary market prices described in Figure 1. Working with the estimates of my structural model I can evaluate the shading in a broader sample (not limited to those with trade reporting requirements to TRACE), and using a more direct measure of bidders' willingness to pay.

6.3. **First-stage behavior.** Analytic solutions for the first-stage optimal strategies are unavailable and numerical solution of these strategies would require calculating the expected profits in stage 2 for every own submission and set of opponents' submissions in round 1 conditional on the vector of private information. Instead of solving these strategies I simply present the pattern of choices observed. I examine the correlations using the estimated private information together with the raw data on first stage submissions. Even post-estimation we do not pin down the private information for a particular individual and so this calculation is done by integrating over the set of possible draws in $\mathcal{M}_i$.

TABLE 10. First Stage

This table presents the correlation of the private information with the choice of y. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

	y_i
n	0.009***
	(0.003)
s_1	0.0005
	(0.003)
s_2	-0.034
	(0.041)
$\eta - R$	-60.566***
	(0.353)
constant	6.278***
	(0.044)

The regression estimates show that the size of the initial submission is positively correlated with the CDS position in the full sample. This is expected based on the auction rules. The expectation of the common value component relative to the opponents expectations also seems to play a key role: bidders with high signals about this component submit substantially smaller physical settlement requests (sell fewer bonds).

6.4. **Evaluating Assumptions.** In setting up the model I made four important assumptions. First, I assumed that bidders bid competitively. The presence of a post-auction resale market allows us to test this assumption. In the presence of collusion the values estimated from a competitive model would fall below the true values. Therefore, if we estimate a competitive model in the presence of collusion, we would expect to see bidders

buying bonds in the post-auction secondary market at prices above the highest estimated values. Section B.3 provides evidence that the observed post-auction trades occur at prices close to the estimated values from the competitive model.

Second, I assumed that bidders truthfully report their initial price quotes. I evaluate this assumption in Section B.4. First, I compare the expected price change a bidder could achieve by manipulating their IMM quote with the size of a fine and show that the fine is much larger than the expected benefit of the small price change. Second, I examine the correlation between a bidder's own quote and their quantity submission: if a bidder is using the quote to manipulate the outcome, these should be positively correlated. Instead, I find a small negative correlation. Third, I assumed that conditioning on the IMM is sufficient to capture all the relevant across-auction heterogeneity in the bonds. To confirm that this assumption is reasonable, Section B.5 presents a set of regressions showing that bond traits have no explanatory power for bids conditional on the IMM and open interest.

The last important assumption is the linearity of the marginal value curve. To test the linearity assumption I first show the R-squared from the within-bidder fit for bidders with more than 3 steps is high, and then show that the estimated positions and value curves are highly correlated with the estimates from re-estimating the model with the inclusion of a quadratic term. Results are presented in Section B.6.

6.5. **Decomposition.** In this section I perform a decomposition to understand the role of the various strategic channels that produce the observed bidding behavior. To do this, I present a partial equilibrium exercise which eliminates various strategic impacts, and allows each individual bidder to re-optimize their bids.

The dynamics in the current two-stage format result in three main features. The first is learning based on the NOI after the first round. Learning from the NOI can be decomposed in to two different parts: learning about the total supply that is offered, and learning about opponents' private information, resulting from the fact that the NOI is the result is constructed from the set of endogenous quantity commitments of all opponents. The second force in the current auction format is the second round quantity constraints: if the second stage were a double auction, in the relevant price range some bidders might like to submit bids supplying the good and some submit bids to demand it, however the current format restricts bidders' possible expressions to either supply or demand (depending on the NOI). This results in the exclusion of some bidders who are no longer able to express their preferences. The third, is the position-reduction effect. In the current format when a bidder commits to y_i in the first round, it effectively reduces the baseline number of units for which they will pay or receive the final price. This position reduction reduces the asymmetry across bidders.

I consider three separate experiments to capture these effects. In the first, I experiment I shut down all learning. This experiment allows bidders to adjust their bids as if the total supply were not announced. I calculate the bidders expectations by constructing simulated residual supply curves where both opposing bids and quantities are drawn randomly. In the second experiment I eliminate uncertainty about the total offered quantity. In this exercise the expected bids by opponents are unchanged from the case without

announcement, but the bidder now knows the total quantity up for sale with certainty. This illustrates the role of pure aggregate uncertainty.[50] In the final exercise I replace $(n-y)$ by n (set $y=0$) and recalculate the bidder's optimal bid.

Unfortunately this decomposition does not allow us to separate the role of participation constraints because the main effect operates by making the first-stage commitments bind. Without the constraint on bidding in stage 2 the first stage would simply be cheap talk. Relaxing these constraints would therefore lead to different first stage choices, which cannot be captured in the decomposition exercise.[51] In equilibrium the game without constraints in the second stage would be identical to the double auction.[52]

The results of these exercises are presented in Table 11 for changes in the price level of the bid made for 10 percent of the total quantity offered. Results for 50 percent and 90 percent are similar in all cases except the experiment eliminating the NOI announcement, where the changes are smaller (0.168 and 0.054, units respectively).

The results of the decomposition suggest that the announcement of the open interest has a net pro-competitive effect. The game with no announcement of the NOI features substantially more uncertainty for bidders about the location and shape of the residual supply curves. Bidders no longer receive information about the levels of their opponents signals or about the size of the aggregate mismatch between supply and demand. When subjected to this uncertainty, bidders tend to increase their bid shading. When the same exercise is performed with a fixed total quantity the bidders respond with less bid shading. The only difference between this case and the previous exercise is that bidders face less uncertainty on the location of the residual supply curve. This suggests that the aggregate uncertainty on the level of the residual supply curve, rather than the loss of information about opponents' signals, plays the dominant role in producing less aggressive bidding without the quantity announcement. The decomposition does not allow us to separate the part of the response due to learning about opponents from anticipating their strategic responses to variation in the quantity level. Taken together, when the bidder no longer gains information about the opponents values from the open interest and does not anticipate strategic responses of the opponents to the total quantity, they submit more aggressive bids.

Finally, without any position reduction bids are slightly lower. The effect is somewhat small, with a change of only 0.233 cents per dollar. However, the position reduction effect

[50]This does not capture the pure effect of learning that arises do the endogeneity of the quantity for sale. That effect should account for changes in opposing bids due directly to the quantity level. Therefore, the magnitude of that effect cannot be calculated using a decomposition within a partial equilibrium setup but requires full solution of a new equilibrium in the game.

[51]If you calculate only the change in bids at the second stage, holding fixed the distribution of open interests and opposing bids this only allows the constrained bidders to adjust their response. However, in equilibrium the other bidders would respond to this adjustment and open interests would no longer play a role in second stage bidding.

[52]However, if the first round commitment game is maintained, there may be many equilibria where bidders make some initial commitment which they later offset by shifting their second stage bids.

TABLE 11. Decompositions

This table presents the average difference (New-Original) in the bidders price bid for 10 percent of the total quantity offered in the baseline auction when bidders re-optimize eliminating each of the channels, holding fixed the distribution of opposing bids.

New-Original bids	Mean
No NOI announcement	-1.424
Direct Quantity Effect	1.789
Position Reduction	-0.233

is only relevant for the subset of bidders that submit non-zero first stage requests (43 percent in the data).

7. Counterfactual

The first counterfactual I consider is the change from the current two-stage auction to a double auction format, proposed by Du and Zhu (2017). A major challenge for the CDS auction mechanism is that the final clearing price establishes both the CDS cash settlement amounts, and serves as a price for the exchange of bonds. Because dealers tend to be net owners of CDS, the cash settlement feature provides dealers with a coordinated incentive to manipulate their bids. As a second counterfactual I maintain the double auction design but experiment with different functions linking the cash settlement payout to the bond exchange price, adding flexibility to the current identity mapping.[53]

An important property in establishing the result that the double auction performs better than the current format, is the requirement that the CDS positions of the dealers are net zero (Du and Zhu (2017)). The estimation results (and the raw data explored in Paulos et al. (2019)) suggest that dealers are net buyers of CDS. This introduces an important price bias, as it means that there will be more shading on the demand side of the market than the supply side — which will tend to push prices down. In the model of Du and Zhu (2017) with continuous supply/demand curves the bias is related to $\sum(\frac{1}{(\mathcal{I}_d-1)})n_i$. As you increase the average n_i this term increases, increasing the downward bias on auction prices. Given the estimated CDS positions, the average n_i implies a price bias of roughly 5.58 cents on the dollar under the double auction format. In the rest of the section I examine this bias when extending the model to allow for step-function bidding and private draws for the slope of the marginal value curve.

The computation of equilibrium in multiunit auction models with step-function bidding has so far been an intractable problem. The challenge arises as equilibrium bid strategies map high-dimensional values $v(q)$ into high-dimensional sets of K_i price-quantity pairs.

[53]Further improvements may be possible if the auctioneer can use regulation to require the CDS position to be reported truthfully, for example via a CCP, rather than relying on the mechanism to illicit reports of this quantity from the dealers. I do not consider any counterfactual changes of this form, as such a mechanism would no longer be solving the same problem and is likely to be able to generate substantial improvements in performance.

Furthermore, these strategies may be highly nonlinear and little is known about their properties. This makes standard methods for numerical computation of these functions infeasible. In a companion paper, Richert (2021), I develop a method to directly compute the equilibrium in these settings without imposing parametric forms on the strategy functions. My method allows for the computation of the entire set of equilibria associated with estimates from a set-identified model in a single run, which is helpful in the multi-unit auction context, as most models of these auctions are only set-identified. In the companion paper I discuss the general properties of the method and provide examples to demonstrate its performance. In Section 7.1, I explain how the method can be applied in the current setting. Section 7.2 provides the solution details. Results are presented in Section 7.3.

7.1. **Counterfactual Solution Method.** I propose to numerically solve for the equilibrium distribution of bids taking as given the distribution of values estimated from the data and the set of equations characterizing equilibrium behavior. To numerically solve for this distribution, I search for the set of bid-distributions for which the distribution of types (eg. private values) that rationalizes these bids in a Bayes-Nash equilibrium matches the known primitive distribution of types. The search proceeds in four steps: (i) guess a bid distribution, (ii) use the model equilibrium constraints to map the bids to values, (iii) check: is the implied distribution of values the same as the known value distribution, (iv) if not: update the guess of the bid distribution and repeat steps (i)-(iii). This procedure can be formalized as the solution to a problem that is very similar (and in some cases equivalent) to indirect inference. The approach differs from existing numerical methods in important ways: (i) for each parameter vector the equilibrium constraints are satisfied exactly, (ii) it imposes parameteric restrictions on the simulation process for bids, thereby allowing the strategy functions to be unrestricted, and (iii) it adopts a different criterion to measure fit. These features allow the entire set of counterfactuals consistent with a set identified model to be computed in a single run of the algorithm.

7.2. **Solution Details.** In this section I discuss the choice of parametrization and criterion function (to measure the distance between implied and true values) in the counterfactual double auction game. These details are not required to understand the results presented in Section 7.3 and can therefore be skipped if the reader so chooses.

In this setting, a "bid distribution" is the joint distribution of prices, quantities, and steps: $G_{\boldsymbol{B},\boldsymbol{Q}|K}(b_{i1},\ldots,b_{ik_i},q_{i1},\ldots,q_{ik_i}|k_i)\pi_K(k_i)$, where π_K is the distribution of steps. For this Monte Carlo exercise I restrict the strategy space to $\bar{K}=8$.[54] The bid distribution is described using sixteen parameters. Given this parametrization of the bid distribution, for any value of $\boldsymbol{\gamma}$, I can simulate the distribution of residual supply curves. Then I can back out the implied private value distributions using the system of first order conditions (equation 3 derived in Section 3).

I parameterize the bid distribution by describing the distribution of quantity levels and price increments. I use a simulated set of 1000 bidders. For each bidder, I draw K_i

[54]Results are robust to alternative choices of $\bar{K}$.

increment pairs, where K_i is sampled uniformly on the support $[0, 1, 2, 3, ..., 8]$. For each bidder I then draw a set of (e_k, f_k), which describe the price change and quantity level from a baseline at each of the K_i steps. With these in hand we have $b_k = \sum_{k'=1}^{k} e_{k-k'} + \bar{\gamma}_p$ and $q_k = \sum_{k'=1}^{k} f_k + \bar{\gamma}_q$, where $\bar{\gamma}_p$ and $\bar{\gamma}_q$ are parameters that determine price and quantity level shifts that apply to all bidders. In this specification I allow the (e_k, q_k, q_{k-1}) to be correlated. I parametrize the marginal distributions of e_1 and f_1 using 4-parameter cubic B-splines, $G_{E_1}(\cdot; \boldsymbol{\gamma}_e)$ and $G_{F_1}(\cdot; \boldsymbol{\gamma}_q)$, characterized by parameter vectors $\boldsymbol{\gamma}_{e_1}$ and $\boldsymbol{\gamma}_{f_1}$ while the marginal distribution of $G_{F_k}(\cdot; \boldsymbol{\gamma}_f)$ for $k \in (2, ..\bar{K})$ as a beta distribution with parameter vector $\boldsymbol{\gamma}_f$ and $G_{E_k}(\cdot; \boldsymbol{\gamma}_e)$ as an beta distribution with parameter vector $\boldsymbol{\gamma}_e$for $k \in (2, ..\bar{K})$. I model the correlation structure as a Gaussian copula $\mathcal{C}[\cdot, \cdot; \boldsymbol{\gamma}_c]$, where $\boldsymbol{\gamma}_{c_2}$ is a 2×2 correlation matrix with elements ρ_{eq} and $\boldsymbol{\gamma}_{c_3}$ is a 3×3 correlation matrix ρ_q, ρ_{eq}, and the third correlation is restricted to be ρ_q, ρ_{eq} which gives conditional independence between e_k and q_{k-1} given q_k:

$$G_{\boldsymbol{E},\boldsymbol{Q}|K}(e_1, \ldots, e_5, f_1, \ldots, f_5 | K_i) = \mathcal{C}\left[G_{E_1}(e_1; \boldsymbol{\gamma}_e), G_{F_1}(f_1; \boldsymbol{\gamma}_f); \boldsymbol{\gamma}_{c_2}\right] \times \prod_{k=2}^{K_i} \mathcal{C}_3\left[G_{E_k}(e_k; \boldsymbol{\gamma}_e), G_{F_k}(f_k; \boldsymbol{\gamma}_q), G_{F_{k-1}}(f_{k-1}; \boldsymbol{\gamma}_q); \boldsymbol{\gamma}_{c_3}\right].$$

Finally, to account for the fact that the probability of K_i steps is not uniform, for each simulated bidder I calculate a weight that reflects the probability of appearing with K_i-steps. To specify this I assume that the probability of putting each additional step is Poisson, with parameter γ_n. For notational convenience I collect all the relevant parameters into a single vector $\boldsymbol{\gamma} = [\bar{\gamma}_q, \bar{\gamma}_p, \boldsymbol{\gamma}_e, \boldsymbol{\gamma}_{e_1}, \boldsymbol{\gamma}_{f_1}, \boldsymbol{\gamma}_f, \boldsymbol{\gamma}_{c_2}, \boldsymbol{\gamma}_{c_3}, \gamma_n]$.

For the criterion function I match the distance between the CDFs of the marginal distributions at a grid of points $\boldsymbol{L}_l$ for $l \in \{s_1, s_2, n\}$ defined by $F_l^{-1}(\alpha) = \boldsymbol{L}_l$ for $\alpha = (0.01, 0.02, 0.03, ..., 0.99, 1)$ and fit the element-wise squared distance between the off-diagonal elements in the matrices of estimated correlations. I base the calculation of these bid distributions off the set of bidders who submit three or more steps. For each bidder I calculate the selection probabilities by fitting a function $h_p(s_1, s_2, n, K)$ using the observed probabilities of s_1, s_2, n under $K = 3, 4, .., \bar{K}$ and then extrapolating this for $K = 1, 2$.[55] Finally, the system of FOC directly may be ill-conditioned for some simulated bids under some parameters of the bid distribution, resulting in large jumps of the criterion function.[56] To improve this, I integrate over a grid of CDS positions n and at each grid point solve the best s_1, s_2 and assigning a relative likelihood to each n by assuming the errors at each of the grid points are normally distributed.[57] Note that as the grid gets fine, in n and the assumed error variance in the normal density, playing a role akin

[55]I also attempted this using h to fit the change in marginal benefit which can then be combined with the estimated marginal cost distribution to calculate the probability of submitting at least three steps. However the high slope of the CDF of marginal costs means that this method is quite sensitive to small changes in the parameters and this makes the criterion function difficult to optimize.

[56]The poor numerical properties come from the value of the third column of the matrix A which for some bid distributions can be close to zero, or constant.

[57]The points n are now simply n as there is no first round submission so $y = 0$.

to a bandwidth in a kernel, goes to zero this is equivalent to the direct solution procedure. The smoothing steps helps to provide good numerical performance of the search and optimization procedures.

To perform the search procedure I first do a grid search and then run local minimization algorithms from the best 1 percent of points. For the grid search I start with a grid of 24.3 million feasible points. I compute the value of the criterion function at one thousand of these points and then fit a neural net approximation (with number of points already evaluated/40 layers) to the criterion function. Using this approximation I then predict the value at all the grid points. Next, I select the 200 grid points that have not yet been evaluated with the best predicted criterion value and compute the criterion at these points. I then refit the approximation to the criterion function and repeat for another 200 grid points. I perform this procedure until the criterion has been computed at 12000 grid points. I then take the best 1 percent of points as starting values for an adaptive mesh search. The solutions from each local adaptive mesh search are then used as starting values for a gradient based sequential quadratic programming solver. I collect the best solution vectors from this to create the solution set.

7.3. **Results.** Two benchmarks provide a useful baseline for comparison to the counterfactual results. First, the counterfactual of truthful bidding in these auctions. This would be the result if there were no strategic bid shading. Second, the outcome under the current non-standard auction rules.[58]

When computing the counterfactual equilibrium I fix the common signal at its median $R = 32.375$. I expect similar shading across different levels of this conditioning variable. To predict the amount of shading in the current format at this level, I predict the gap from the IMM to the auction price. Across auctions this has an average of -0.1082 and the size of the gap is independent of the IMM level. If I separate auctions with $NOI > 0$ and $NOI < 0$, the average gaps from the IMM to auction price are -1.5043 and 2.6875, respectively. This implies that for $R = 32.375$, auction prices in the current format should be 32.27 while under truthful bidding they should be $32.375 + (4 - 0.1082) = 36.27$.

The main policy counterfactual is a change to a double auction format. The results of this exercise suggest that the double auction could increase the price in the auctions to 33.92 from 32.27 today. This represents a 41.25 percent decrease in the level of bid shading from the current format. However, the double auction still results in substantial shading, as the market clearing price is on average 2.35 cents on the dollar below the value under truthful bidding.

[58]It is not possible to solve the equilibrium of the current auction format and so I compare outcomes to the data. In the fully non-parametric case, the outcome in the data would be equivalent to the model equilibrium predictions. The main parametric restriction is the linear form of values; to show that this does not drive results I calculate the optimal bid for each bidder who submitted three or more steps (allowing me to pin down their $s_1, s_2, (n - y)$) imposing this linear form on values. I then compare the calculated optimal bid to their observed bid. The resulting bids are very similar: in 95 percent of cases the change in the expected clearing price conditional on the bid being made is less than 1e-10 and so it seems unlikely that this is what drives the results.

TABLE 12. Change in Auction Format

	Data				Double Auction	
	$NOI > 0$	$NOI < 0$	All	Truthful Bidding	BL	No CDS
Auction Price	30.87	35.06	32.27	36.27	33.92	35.86

The counterfactual change to a double auction reduces the risk faced by investors in two ways. First, it directly reduces the auction outcome risk. Outcome risk is generated by the fact that the bias in any given auction is unpredictable and can be measured using $Var(p^{auc} - E[p^{auc}|R])$. The current auction format has a standard deviation in these outcomes of 5.56 cents/dollar (or 3.37 when outliers are ommitted). The counterfactual double auction reduces this substantially, to 2.32 cents/dollar. The second source of risk is the risk directly generated by the price bias. Plots illustrating the role of this bias are provided in Appendix D. Because the bias is a fixed cents/dollar rather than a percentage of the final recovery price, investors and because the recovery amount is not known before the credit event occurs, investors cannot simply adjust their holdings to offset the role of the pricing bias. If investors adjusted their positions to account for the expected level of recoveries, they would be underpaid when recoveries are low and overpaid when high. Given the large bias in the current auction format this leads to an additional risk to investors with a standard deviation of 2.05 cents/dollar, which is reduced under the counterfactual auction format to 0.88 cents/dollar.

A major challenge for the CDS auction mechanism is that the final clearing price jointly determines (i) the CDS cash settlement amounts, and (ii) the price for bonds exchanged. Because dealers tend to hold net positions on the same side of the market, the cash settlement feature provides them with a coordinated incentive to manipulate their bids. In the next exercise I isolate the bias introduced from dealers' CDS positions. To do this, I compute a counterfactual double auction where bidders do not consider the impact of the auction price on their CDS payouts when choosing their bids.[59] The expected clearing price from this exercise is close to the price under truthful bidding, suggesting that the price bias is largely driven by dealers' CDS positions.

To evaluate the efficiency of the auction format I compare the expected surplus of bidders under the current and double auction designs to the surplus they would obtain if the bonds were assigned under the truthful bidding benchmark. The current design achieves only 15% of the possible surplus. The double auction improves on this substantially achieving 22% of the possible surplus. However, both the current and double auction designs achieve fairly inefficient allocations, as in both cases the allocations are heavily influenced by the positions of participants in CDS contracts which are irrelevant under the efficient benchmark.

[59]Because the CDS positions play no role in this exercise, the criterion function for the computation only evaluates the marginal distributions of signals.

The extra risk in these contracts reduces the usefulness of the insurance contracts and represents a loss of welfare from a contract with full insurance. This reduced ability for investors to insure themselves could increase the costs of holding bonds, and could reduce the gains to firms of having CDS written on their debt. To understand the welfare impact, I first calculate from a world of no insurance the percentage of the way towards full insurance that is covered by the current contracts. I then make the same calculation using the insurance provided when the current auction is replaced with a double auction. The results suggest the current contract achieves 81-87% of the total gains well the double auction would achieve 93-95%. To benchmark these gains to the real outcomes from these contracts, I use the estimate from Danis and Gamba (2018) which suggests that the current CDS contracts cause an increase in firm value of 2.9% when they are introduced on a firm. Notably, they find the presence of these contracts on being written on a firm increases the firms investment and leverage.[60] This would suggest that the replacing the current auction rules with a double auction would increase the firm value by 3.18-3.35% instead of 2.9%, or an additional increase in firm value of 0.23-0.45% representing a huge missed gain.[61]

In appendix E I evaluate the sensitivity of these results to changes in the CDS or bond positions of dealers in response to the change in auction format. To do this I define bounds on the changes in the joint distribution and recompute the set of counterfactual outcomes consistent with these bounds. Even allowing for these changes, the double auction should not perform substantially worse than the current auction format.

8. Conclusion

In this paper I develop and estimate a structural model describing bidding behavior in credit event auctions. The current auctions have two stages with bidders providing initial quantity commitments to buy or sell fixed quantities of bonds and then clearing the excess supply or demand using a uniform price auction. To model these auctions, I extend models of bidding in multiunit auctions to handle the initial positions. I then show how bidding data can be used to identify both the private values and CDS positions of dealers without placing parametric restrictions on the shape of dealers marginal value functions. Given this, I estimate the private information from bidding behavior in CDS auctions and use the estimates to perform a decomposition exercise of the importance of a set of strategic channels. Then I apply a novel computational tool which I develop in Richert (2021) to directly solve for the counterfactual equilibrium in multiunit auction games to study the outcome of a change to a uniform price auction format.

I find that the current design results in substantial market power for the dealers and as most dealers hold CDS, prices that are substantially lower then their true willingness

[60]This estimate may be best thought of as an upper bound, some earlier work including Ashcraft and Santos (2009) fail to find evidence of an impact from CDS on bond markets though other recent work including Bilan and Gündüz (2022) also finds evidence of a positive link.

[61]Because the default probability of the firm is unknown I compute the risk reduction implied by a grid of default probabilities and report the maximum and minimum.

to pay. This is largely driven by the fact that the auction determines both allocation and payments for bonds and the cash settlement amount for CDS positions. Because of this, dealers use their bids both to express their value for bonds and to influence their payments owed on their CDS positions. While the current format improves competition by providing information between rounds that reduces the aggregate uncertainty, and the first stage commitments reduce bidders' CDS position exposed to the auction price, the constraints on participation stop many dealers from expressing their valuations. Switching to a double auction reduces the amount of shading by 41.25 percent from the current level. This leaves a still substantial amount of bid shading, largely because dealers are mostly buyers of CDS and so have an incentive to push auction prices down. I then show that without the incentive for bidders to shade their bids to profit from their CDS positions, the double auction results in almost no price bias.

The size of bid shading in credit event auctions is large and the prices are substantially below the dealers willingness to pay for bonds. These effects mean that when a credit event occurs the sellers of insurance are responsible for making payments of hundreds of millions beyond what they would owe under fair insurance and the amount of payments owed varies substantially due to the information rents extracted by large dealer banks. These frictions reduce the insurance provided by these contracts and expose holders to additional risk. Given the size of this additional risk, this has important implications for the functioning of the CDS market. A benchmark calculation suggests that the auction design means we are missing out on important increases in firm value for the large firms covered by these insurance contracts leading to a substantial cost for the real economy.

References

Allaz, B. and J.-L. Vila (1993). Cournot competition, forward markets and efficiency. *Journal of Economic Theory 59*(1), 1–16.

Armantier, O., J.-P. Florens, and J.-F. Richard (2008). Approximation of Nash equilibria in Bayesian games. *Journal of Applied Econometrics 23*(7), 965–981.

Ashcraft, A. B. and J. A. Santos (2009). Has the cds market lowered the cost of corporate debt? *Journal of Monetary Economics 56*(4), 514–523.

Augustin, P., M. G. Subrahmanyam, D. Y. Tang, and S. Q. Wang (2014). Credit default swaps: A survey. *Foundations and Trends(R) in Finance 9*(1-2), 1–196.

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014). Demand reduction and inefficiency in multi-unit auctions. *Review of Economic Studies 81*(4), 1366–1400.

Bao, J., J. Pan, and J. Wang (2011). The illiquidity of corporate bonds. *The Journal of Finance 66*(3), 911–946.

Bertrand, A., I. Van Keilegom, and C. Legrand (2019). Flexible parametric approach to classical measurement error variance estimation without auxiliary data. *Biometrics 75*(1), 297–307.

Bessembinder, H., C. Spatt, and K. Venkataraman (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis 55*(1), 1–45.

Bilan, A. and Y. Gündüz (2022). CDS market structure and bond spreads. Technical report.

Bonaldi, P. (2017). Motives and consequences of LIBOR strategic reporting: How much can we learn from banks' self reported borrowing rates. Technical report, Carnegie Mellon Tepper Schook of Business.

Chernov, M., A. S. Gorbenko, and I. Makarov (2013). CDS Auctions. *Review of Financial Studies 26*(3), 768–805.

Conley, T. G. and F. Decarolis (2016). Detecting bidders groups in collusive auctions. *American Economic Journal: Microeconomics 8*(2), 1–38.

Coudert, V. and M. Gex. (2010). The credit default swap market and the settlement of large defaults. *Économie internationale*, 91–120.

Danis, A. and A. Gamba (2018). The real effects of credit default swaps. *Journal of Financial Economics 127*(1), 51–76.

Das, S., M. Kalimipalli, and S. Nayak (2014). Did cds trading improve the market for corporate bonds? *Journal of Financial Economics 111*(2), 495–525.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266–284.

Dick-Nielsen, J. (2009). Liquidity biases in TRACE. *The Journal of Fixed Income 19*(2), 43–55.

Du, S. and H. Zhu (2017). Are CDS Auctions biased and inefficient? *Journal of Finance*, 47.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Feldhütter, P., E. Hotchkiss, and O. Karakaş (2016). The value of creditor control in corporate bonds. *Journal of Financial Economics 121*(1), 1–27.

Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica 68*(3), 525–574.

Gupta, S. and R. K. Sundaram (2012). CDS Auctions and Informative Biases in CDS Recovery Rates. *SSRN Electronic Journal*.

Gupta, S. and R. K. Sundaram (2015, November). Inventory effects, the winner's curse, and bid shading in Credit Default Swap auction outcomes. *The Journal of Derivatives 23*(2), 36–52.

Hansch, O., N. Y. Naik, and S. Viswanathan (1998). Do inventories matter in dealership markets? Evidence from the London stock exchange. *The Journal of Finance 53*(5), 1623–1656.

Hickman, B., T. Hubbard, and E. Richert (2021). Testing independent private values versus alternative models in auction models. *mimeo*.

Hortaçsu, A. and J. Kastl (2012). Valuing dealers' informational advantage: A study of Canadian Treasury auctions. *Econometrica 80*(6), 2511–2542.

Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish Treasury auction market. *Journal of Political Economy 118*(5), 833 – 865.

Hugonnier, J., B. Lester, and P.-O. Weill (2019). Frictional intermediation in over-the-counter markets. *The Review of Economic Studies 87*(3), 1432–1469.

Ito, K. and M. Reguant (2016, July). Sequential Markets, Market Power and Arbitrage. *American Economic Review 106*(7), 1921–57.

Judd, K. (1998). *Numerical Methods in Economics*. MIT Press.

Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *The Review of Economic Studies 78*(3), 974–1014.

Kastl, J. (2012). On the properties of equilibria in private value divisible good auctions with constrained bidding. *Journal of Mathematical Economics 48*(6), 339–352.

Kastl, J. (2020). Auctions in financial markets. *International Journal of Industrial Organization 70*, 102559.

Klemperer, P. D. and M. A. Meyer (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica 57*(6), 1243–1277.

Kremer, I. and K. G. Nyborg (2004). Divisible-good auctions: The role of allocation rules. *The RAND Journal of Economics 35*(1), 147–159.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *The Review of Economic Studies 56*(3), 317–355.

Laksa, S., D. Marszalec, and A. Teytelboym (2018). Epic fail: How below-bid pricing backfires in multiunit auctions. CIRJE F-Series CIRJE-F-1096, CIRJE, Faculty of Economics, University of Tokyo.

Li, D. and N. Schurhoff (2019). Dealer networks. *The Journal of Finance 74*(1), 91–144.

Paddrik, M. and S. Tompaidis (2019). Market-making costs and liquidity: Evidence from cds markets. Working Papers 19-01, Office of Financial Research, US Department of the Treasury.

Paulos, E., B. Sultanum, and E. Tobin (2019). CDS auctions: An overview. *Economic Quarterly*.

Peivandi, A. (2015). Participation and unbiased pricing in CDS settlement mechanisms.

Richert, E. (2021). Indirect inference techniques as an alternative to direct computation of equilibrium bid distributions. *mimeo*.

Rutledge, W. (2009). Observations on management of recent credit default swap credit events. Report, Senior Suprvisors Group.

Slive, J., J. Witmer, and E. Woodman (2012). Liquidity and Central Clearing: Evidence from the CDS Market. Staff Working Papers 12-38, Bank of Canada.

Vives, X. (2011). Strategic supply function competition with private information. *Econometrica 79*(6), 1919–1966.

Wilson, R. (1979). Auctions of Shares. *The Quarterly Journal of Economics 93*(4), 675.

Zhang, A. L. (2022). Competition and manipulation in derivative contract markets. *Journal of Financial Economics 144*(2), 396–413.

Appendix A. Equilibrium Existence

To understand the role of the restricted strategy sets it is useful to compare the results to the unrestricted case from Wilson (1979) in the IPV case. Using calculus of variations gives $v(q,s) = b-(q+n-y)\frac{H_q(b,q|s\Omega)}{H_b(b,q|s\Omega)}$, where H represents the probability that the residual supply is less than or equal to the quantity q at price b. Using this together with Proposition 4 from Kastl (2012) implies that as K goes to infinity, any restricted equilibrium approaches this solution and these empirical FOC are valid for inference conditional on an equilibrium existing.[62] Although the existence of an equilibrium in the uniform price auction with restricted strategy sets is an open question, in the standard setting Kastl (2012) proves the existence of an epsilon equilibrium. This argument cannot be applied to the credit event auction setting, because the proof makes use of the separability between the benefit of winning and the price paid, to argue that if a bidder is unrestricted in number of steps they will not bid above their value. This separability property does not apply in credit event auctions, as a bidder may be better off bidding above their value in order to impact the clearing price of their existing CDS position. To guarantee that an equilibrium exists, in the uniform price multiunit auction game, I follow the suggestion of Kastl (2011) and impose that there exists a fine discrete grid of price levels. This is the case in practice, as bidders can only express their prices to the nearest 1/8th of a cent. In this case, Kastl (2011) argues that the FOC for the quantity choice are still valid, and an equilibrium is guaranteed to exist (at least in mixed strategies) as it is a finite game.

Appendix B. Additional Tables

B.1. **Additional Summary Statistics.** Table A.1 presents some additional statistics describing the bidding behavior of different bidders. The sample is loosely divided into bidders that are involved regularly (the 9-10 global dealers) and the less frequent regional participants whose participation varies with the frequency of defaults in that location in the sample period. The table also shows there is considerable variation within participant in the direction of their initial quantity commitments, suggesting that while the dealers are most often holders of insurance, they are still net sellers in some cases. It also appears to be quite common that bidders submit only 1 step in the second stage. These bidders do not actively bid in the second stage as 1 step is carried-over from the initial price quotes.

Figure A.1 shows the price realized in auctions depending on which type of credit event caused the auction. Even across types of credit events, there is a large amount of heterogeneity in the remaining value of the firms. However in regressions, there is no evidence that the type of event is correlated with the gaps between the pre-auction price, auction price, and post-auction resale prices.

B.2. **Multiplicative Form.** If the common value component of the bond entered multiplicatively with bidders' own private values in the model then the dispersion of private

[62]The proposition requires randomness in the quantity being sold which is announced in this game. However, the carried over amounts from the first stage price quotes effectively lead to a random (predetermined and non-strategically linked) residual quantity at any price level.

TABLE A.1. Auction Description

The following table presents summary statistics for participation of the bidders. Each number is a count of the number of auctions in which the bidder participated, submitted a positive or negative first round quantity commitment or used each number of steps in their second stage bids.

Bidder	Participated	$y_i>0$	$y_i<0$	1 step	2 steps	3 steps	4 steps	5+ steps
1	1	1	0	0	0	0	1	0
2	123	40	25	36	28	19	8	24
3	175	51	28	57	25	8	12	56
4	4	2	1	1	2	1	0	0
5	135	42	18	81	17	8	8	17
6	169	35	29	68	21	13	8	51
7	171	43	23	64	42	18	8	27
8	12	8	0	9	1	0	1	0
9	179	47	25	38	46	17	17	50
10	47	10	17	35	6	1	0	4
11	1	1	0	1	0	0	0	0
12	183	55	24	51	40	26	11	47
13	4	3	0	1	2	1	0	0
14	68	15	13	35	15	6	1	8
15	2	1	1	2	0	0	0	0
16	5	3	1	5	0	0	0	0
17	1	0	0	0	0	0	0	0
18	180	46	36	53	29	25	14	50
19	78	15	19	50	7	2	2	14
20	102	22	13	50	22	6	7	12
21	96	19	11	83	4	0	1	6
22	1	0	1	0	0	0	0	1
23	180	43	42	72	31	16	13	40
24	2	1	0	2	0	0	0	0
25	128	32	27	67	22	7	4	24

information would be increasing in the level of R. This would mean that for auctions with small R the dispersion in private values matters little, while in auctions with a big R this plays a central role. If this was true we would expect to see the level of information rents, and the gap with the pre-auction price information growing in R. Figure A.2 plots the auction price against the IMM price in each auction. Since there is no evidence for a difference in price gaps across the levels of R, I adopt the additive specification.

B.3. **Collusion Test.** Although the setting features repeated interaction of a small set of participants, it would likely to be difficult to sustain collusion as (i) violation would be difficult to detect (ii) there are usually one or two regional players in each auction who do not frequently participate. Detection is difficult in this setting because bidders may receive orders which they place on behalf of their customers. This means that when deviating

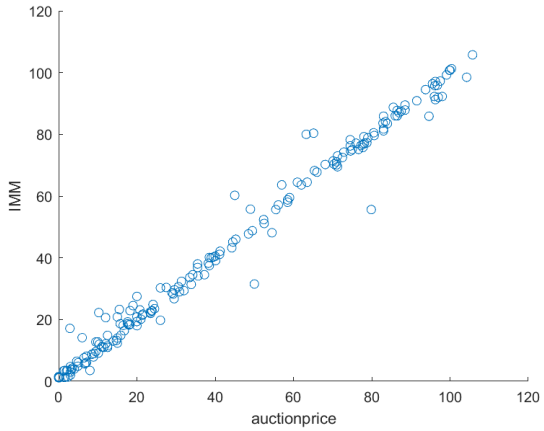
FIGURE A.1. Event Types: Prices

The figure plots the price realized in auctions depending on which type of credit event caused the auction. Even across types of credit events, there is a large amount of heterogeneity in the remaining value of the firms.



FIGURE A.2. Auction price vs NOI

The figure plots the price realized in auctions against the IMM quoted.



from the prescribed collusive behavior, bidders could simply claim to be placing the bid on behalf of a customer. Since this cannot be verified by the other participants this makes detection and deciding when to punish more complicated, making collusion more difficult to sustain. In addition to this, the set of bidders that wants to push the price up and the set that wants to collude to push the price down, (ie. the sets of individuals that all benefit from working together) varies across every auction. Finally if prices were pushed substantially in one direction bidders with large insurance positions on the opposite side of the market would have a strong incentive to deviate from the agreement. These challenges

together with the reduced form evidence mean that collusion in the second stage bidding game would be quite difficult.

The resale opportunity present in the bond market allows for an additional test of the null hypothesis of no collusion. In a model of collusion we would expect bidding behavior similar to that described in Laksa et al. (2018). If the data was generated by collusive bidding, then bidders implied values that rationalize observed bids in a competitive bidding model would be well bellow the true values. If we then saw bidders willing to buy bonds immediately after the auction at higher resale prices this might suggest a violation of the competitive bidding model. In the data, the median value implied is 0.96 cents below the IMM at the expected clearing quantity, however the 65th percentile is the IMM and the 78th is the average markup for the clearing price. At the 90th percentile the value is 5 cents above the IMM and at the 99th it is 45 cents above the IMM. These results seem to be broadly consistent with the observed post-auction behavior and not suggestive of collusion.

B.4. **IMM manipulation.** The average change in price that dealers can expect by manipulating their IMM quote is 0.02 cents. This is small as if you quote a number that is different from others your quote is dropped and since only half of the quotes are used and the average is rounded to the nearest 1/8th of a cent increment after averaging, it is difficult to influence this calculation with a unilateral deviation. At the 95th percentile of expected benefits when integrating over the estimated distribution of possible n and using the distribution of clearing prices in the data, this gives an increase of 4,198 dollars of surplus. The mean cost from quoting off-market is 24,000 so a bidder that is optimizing should be more worried about that effect and quote their best guess of initial price.

Given the incentives to profit from the insurance positions, first-stage quotes should be negatively correlated with the bidders' insurance positions, ie. a bidder who is a net buyer of CDS should quote lower prices in the first round. Table A.4 presents results from regressing the level of individual price quotes on an indicator which takes the value 1 if the nonparametericallly estimated lower bound is greater than zero (column 1) and upper bound is greater than zero (column 2). I use the nonparameteric estimate as these are available for all bidders rather than only those submitting at least 3 steps. While the coefficient has a negative sign, it is not statistically significant in either specification. Table A.3 presents an alternative specification, comparing bidders who commit to sell bonds (and so must be buyers of insurance according to the auction rules). Again the coefficient is negative but not statistically significantly different from zero.

B.5. **Sufficiency of IMM.** Given the strong relationships documented between outcomes and the IMM price, I proposed that the IMM should be considered a sufficient statistic for the auction level heterogeneiety. In this section I show that while there is some evidence that bond traits influence the IMM amounts, there is no evidence that they influence residual bids beyond this point. Therefore, conditioning on the IMM should be sufficient to capture the auction specific differences in bonds.

TABLE A.2. Position and price quotes

The following table presents results from regressing the level of individual price quotes on an indicator which takes the value 1 if the nonparameterically estimated lower bound is greater than zero (column 1) and upper bound is greater than zero (column 2). I use the nonparameteric estimate as these are available for all bidders rather than only those submitting at least 3 steps. In all cases I control for the baseline expected recovery value using the final IMM. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Variable	IMM Submission	IMM Submission
Auction IMM	0.9927***	0.9929***
	(0.0014)	(0.0014)
CDS buyer	-0.1346	-0.1202
	(0.094)	(0.1151)
Constant	Yes	Yes

TABLE A.3. Position and price quotes

The following table presents results from regressing the level of individual price quotes-IMM on an indicator which takes the value 1 if the noi submission is positive. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Variable	IMM Submission
CDS buyer	-0.1296
	(0.235)
Constant	Yes

Finally, I check the relationship of the traits of the deliverable bonds to the auction outcomes. I check this relationship both at the bidder-level, regressing the residualized bids on the bond traits in Table A.4 and at the auction level in Table A.5. This shows that the bond traits have some power in explaining the IMM quote that a bidder provides but no power to explain their residualized bid after conditioning on the IMM. The regression at the auction level finds no statistically significant effect of the bond traits. Given these results, I do not include bond traits in the main estimation. These results indicate that once I condition on the initial market price they have no explanatory power.

Section C discusses in detail the incentives involved in the IMM submission decision and considers a calibration exercise to examine the information revealed through the IMM announcement.

B.6. **Linearity Test.** I test the linearity assumption in two ways. First, by using the overidentifying restriction from the subsample of bidders that submit 3 or more bids. In that sample, the median R squared is 0.98 and the mean 0.87. As a second test I estimate the model with a quadratic specification for marginal values. The median change in the estimated CDS position is 0.015 million and even at the 75th percentile the change is only 1.5 million. The estimated positions under the two sets of CDS positions n are also

TABLE A.4. Bond Traits: Bidder Level

Residualized bids from nonparameteric regression on IMM and NOI. Bonds useful in determining IMM submissison but not in bids conditional on IMM common signal and NOI. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Variable	IMM Submission	Residualized bid
Duration	-4.888**	-0.0177
	(2.437)	(0.0318)
Conversion	5.284	-0.003
	(3.394)	(0.003)
Convexity	0.531**	0.059
	(0.266)	(0.065)
Volume	0.0002	0.011
	(0.0031)	(0.0325)
Auction NOI	Yes	No
Constant	Yes	Yes
N	1965	1965

strongly positively correlated. Because of this, I maintain the linear restriction for the primary specification.

APPENDIX C. STAGE 1 PRICE QUOTES

The bidders choice of first stage price quotes is a complex strategic decision. These quotes serve many roles in the auction (i) the quoted price is carried over as a bid for a fixed quantity (usually 2 million dollars) of bonds in the second stage (ii) the quotes are aggregated by taking an average excluding the outliers which is announced to all participants between rounds (iii) the average plus 2 times the spread determines a price cap or floor which stands for bids submitted in the second stage auction (iv) the quotes determine a set of fines for bidders who submit off-market quotes (those that differ substantially from the average).

These many roles mean that bidders' may have incentives to strategically report their quotes from a number of different sources and their strategies are likely to be complex. The data contain some information that indicates the importance of the different channels. The carried over bids are sometimes relevant for clearing the auction, for example most auctions have carried over amounts of 2 million and an average of 11 participants implying 22 million of carried over bids. 74 auctions have a total excess supply/demand less than 22 million and so these bids may play an important role in the final price determination. The price cap binds for only 3 percent of the bids made in the second stage auctions, but in 16 percent of the auctions it plays an important role in determining the price. Fines are given 169 times in the data and have an average level of 32000 dollars.

Assumption 1, imposes that once bidders know the IMM their own private information on the common value component of bond values is no longer relevant (ie. the bidders mostly agree on the common component of the recovery value). This assumption is

TABLE A.5. Bond Traits: Auction Level

Auction price vs imm influence of bond traits. Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

VARIABLES	(1) br_auc_pauc	(2) br_auc_imm
duration	-0.643	-6.901
	(1.495)	(8.161)
conversion	0.692	3.751
	(1.985)	(10.85)
convexity	0.0538	0.587
	(0.166)	(0.910)
volume	3.25e-06	0.000154
	(1.82e-05)	(9.91e-05)
br_auc_noi	-0.00718***	-0.0389***
	(0.00242)	(0.0129)
br_auc_imm	1.004***	
	(0.0571)	
imm2	-5.19e-05	
	(0.000568)	
Constant	0.599	50.05***
	(1.643)	(6.775)
Observations	178	178
R-squared	0.971	0.068

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

critical for the tractability of the empirical exercise. Although bidders may have many reasons to manipulate their price quote their chosen quote is likely to be correlated with their own signal, which influences their expectations about what opposing quotes they will face. If bidders make reports that are correlated with their own signal, the IMM will aggregate the signals from across many bidders. For bidders trying to learn about the mean the IMM which combines information from many draws is likely to be much more informative than the individual bidders' single information. In addition, the complex formulas for the calculation of the IMM make it difficult for a bidder to calculate the role of their own quote in establishing the final value, reducing the value of relying both on their quote and the announced IMM.

To get a better understanding of the updating process I perform a simulation exercise to understand what bidders learn with parameters calibrated to match the quotes in the data. I impose that each bidder learns a signal about the true common value, and also has draws some private benefit from misreporting their value either up or down. This

private benefit draw is a reduced form way to capture the complex benefits that a bidder may receive from manipulating the price quote. For example the benefit may come from the change in profitability from the bidders carried over bids, or the change in expected profits from their influence on the price cap or floor.

For each bidder I calculate their expected impact on the IMM from submitting different price quotes along a grid of possible submissions ranging from 10 cents on the dollar below the true value to 10 cents on the dollar above the true value. In calculating the expected impact I assume that bidders expect to face quotes drawn from the empirical distribution of quotes submitted into auctions with similar post-auction prices to the common value signal $R + \eta$ that the bidder received. I then calculate for each simulated information set, the quote on the grid that maximizes the bidders' expected surplus given by the private marginal benefit of manipulation multiplied by the price impact of their quote less any fines. I will assume that the distribution of private signals and private benefits are both normal and the private benefit is independent of the private signal about the common value for bonds.

Given this structure there are three critical parameters which are unknown, that will determine the amount of information revealed by the IMM announcement. First, the precision of the initial signals about this component, which is governed by the parameter σ_η. Second, the mean and variance of the distribution of private benefits from manipulation. I set a coarse grid in these three parameters and for each grid point solve the choice of initial quotes as described above for 100 randomly drawn private benefits and initial signals. I then compare the distribution of the implied optimal quotes to the distribution of quotes submitted in the data and choose the parameters which minimize the difference between these two distributions. This gives the key inputs to the updating process: a signal variance, a quote variance and a correlation between the signal a bidder receives and their submitted quote.

Given these parameters I can simulate signals η. For each signal I can calculate the optimal submission for that bidder and after repeating for each bidder at a simulated auction, can obtain a resulting IMM. Using this I can calculate the expectations of each bidder $E[R|\eta]$ and $E[R|\eta, p^{IMM}]$ assuming they update according to bayes rule and know all relevant distributions, and that the underlying distribution of R is exactly equal to the post-auction resale price distribution in the data.

Two quantities play an important role in the outcomes. The variance across bidders (within-auction) of $E[R|\eta, p^{IMM}]$ which indicates the remaining role of the signals η and the ratio of this variance to $E[R|\eta]$ which indicates how much the bidders learned. The participating bidders appear to have a fairly precise knowledge of the common value component with initial expectations having an expected variance of 0.6 cents. Once the IMM is announced this disagreement drops dramatically and bidders almost completely agree with each other. The remaining variance under the calibrated parameters is 0.002 cents, which is roughly .3 percent of the variance in initial expectations. In experimenting with the parameter values it appears that even under quite small correlations between quotes and initial signals the IMM quoting mechanism results in expectations that are far

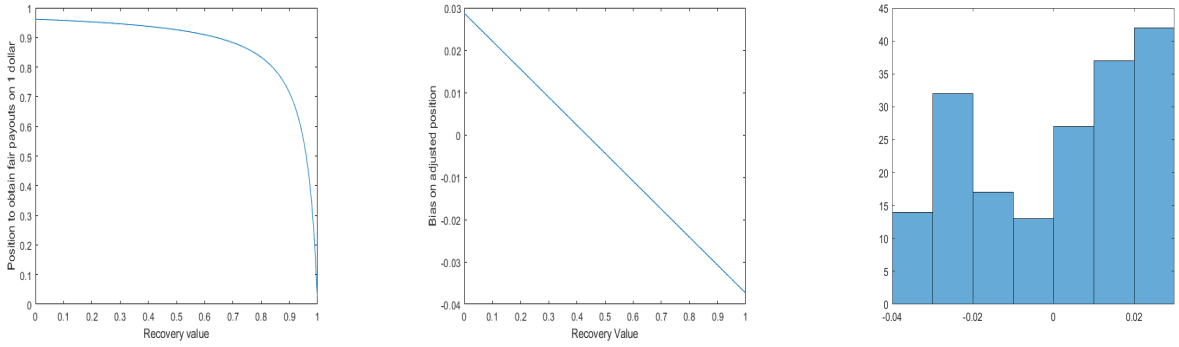
less variable across participants. The small variance in the expectation across different initial signals after learning the initial market quote suggests that heterogeneity in bidders' expectations of the common value post announcement are not likely to play an important role.

While collusion in stage 2 bids may be difficult to sustain it is possible that bidders instead collude on their first stage quotes. However, even under collusion it seems likely that the optimal quote level depends on the initial signals received by bidders and so the level chosen is likely to be highly informative to bidders of opponents signals, reducing their reliance on their own initial signal. For example, when a bidder has a high η, they are willing to buy bonds at higher prices. Making a low quote would decrease the price cap making it more likely that they are constrained, allowing them to purchase fewer bonds at the attractive price. They would therefore want to bargain for a slightly higher quote and the final IMM would reflect this information weighted against the other collusive participants. Therefore it seems likely that even under a collusive regime, the IMM level would substantially reduce the reliance of bidders on their initial signals of the common value. Any bias due to collusion in these first stage quotes is captured in estimation by the function $R(IMM)$, which is parameterized as a cubic B-spline.

Appendix D. Risk Calculation

The following figures illustrate the risk induced by the fact that there is a constant level of expected bias while recovery values are uncertain before the auction. Results are shown for the level of bias under the current auction format and under the counterfactual double auction design.

Figure A.3. Bias-induced Risk

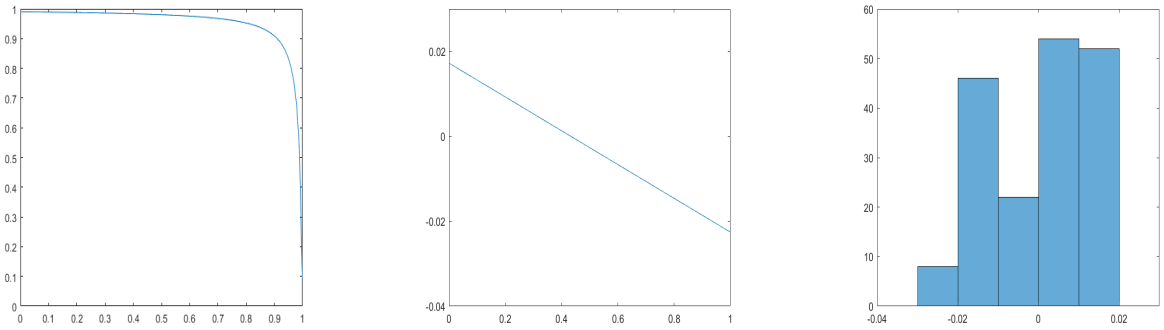


(A) Average Adjustment (B) Bias across Recovery Values (C) Realized Bias

Appendix E. Counterfactual Robustness to Changes in Positions

In the results so far I have assumed that the joint distribution of s_1, s_2, n was a primitive and would remain fixed in the counterfactuals. This assumption seems reasonable given

FIGURE A.4. Bias-induced Risk Counterfactual



(A) Average Adjustment (B) Bias across Recovery Values (C) Realized Bias

the CDS and bond positions are taken on prior to the default event occuring.[63] Therefore, they are likely to be much more reflective of market-making and trading activities by the dealers, their costs of holding bonds and CDS, and their perceptions about the probabilities of default than the expected auction outcomes. However, we may be worried for example that a reduction in the expected surplus at the auction from holding CDS leads dealers to hold a smaller initial position. In this section I discuss the most plausible ways that the joint distribution might be affected by the change in incentives to hold various positions in the counterfactual auction formats. I then develop a set of changes to the value distribution that are plausible and use this to compute a set of bounds for counterfactual equilibria for any joint distribution in the set.

There are two possible changes that one may worry could occur that would affect the joint distribution. The first, are changes to the CDS position caused by shifts in the benefit of holding a particular CDS position for a given marginal value curve and given pre-auction benefit of holding CDS. The second, are changes to the bonds bought/sold before the auction which could shift bidders along the marginal value curve (ie. lead to bidding behavior according to $v(q) = s_1 - s_2\Delta B_i - s_2 q$).

First consider changes in the CDS position. These changes may play an important role through the constraints they impose on a bidders' set of feasible actions. For example, these constraints may prohibit a bidder from obtaining their desired final position in bonds. This incentive is discussed at length in Du and Zhu (2017) and they show that under the current auction format the desire to be unconstrained leads bidders with intermediate levels of pre-auction benefit from holding CDS on both the buy and sell side to hold slightly larger positions. The lack of constraints in the double auction should eliminate this expansion. In the double auction, bidders also no longer have the option of a physical settlement round. Given the concentration of buyers/sellers I still expect the double auction to achieve a downward bias in general on the price, which could provide an

[63]There is also a limited amount of trading (and limited liquidity) that takes place in the lead up (and during) the auction. For the trace-eligible sub-sample of auctions the median trade volume on the auction day is \$6.5M of bonds.

incentive for buyers to increase their positions and sellers to decrease their positions (such that n, rather than $n - y$ is subject to the price bias). These shifts in the distribution will increase price biases in the CDS auctions and so the baseline results may overstate the possible improvement. Because it is likely that most of the position is determined by factors unrelated to the auction, I consider as a reasonable set of bounds, perturbations that allow for an increase of up to +10 percent of each CDS buying bidders existing CDS position and a decrease of 10 percent on the positions of seller dealers.

Given the expected price pressures from cash settlement, bidders expect the bonds traded in the auction to do so at a discount to the market price of bonds in both the current and double auction format. In the baseline change to a double auction there is a slight reduction in the level of the discount for bonds purchased in the auction. This would suggest that bonds purchased in the auction are relatively less attractive and may lead high value bidders to purchase additional bonds before the auction date. This change in positions is expected to lead to less aggressive bidding and lower prices, so the main double auction results may only be an upper bound on the possible set. The bond market is quite illiquid and especially so following default (as documented in Feldhütter et al. (2016)) and so large adjustments of positions will generally be extremely costly. Therefore I examine robustness of the results to a shift in the intercept distribution that is consistent with a shift along the value curve equivalent to a maximum purchase of $1 Million of bonds by high (above median) value bidders prior to the auction.

The results of this exercise suggest the final price after adjustments in position will be in the interval 31.03-34.47. This means that once the position changes are accounted for, the double auction may fail to improve the current design but it appears unlikely to perform substantially worse than the current format.