Nuffield College
UNIVERSITY OF OXFORD

ECONOMICS DISCUSSION PAPERS

Paper No: 2024-W02

Competing Models of the Bank of England's Liquidity Auctions: Truthful Bidding is a Good Approximation?

By Charlotte Grace

March 2024

Competing Models of the Bank of England's Liquidity Auctions:

Truthful Bidding is a Good Approximation

Charlotte Grace

Nuffield College, University of Oxford*

March 24, 2024

Abstract

This paper provides a method for comparing different models of behavior in auctions. It uses data on participants' bids but not their values. I find that a model of truthful bidding (bidding one's true value) outperforms a conventional model (shading one's bids to maximize expected surplus) in the Bank of England's uniform-price liquidity auctions. Truthful bidding is preferable if either (i) bidders are risk averse to a degree that is consistent with that found in the studies closest to my setting, or (ii) the cost of calculating what would otherwise be the optimal strategy exceeds around 5% of bidder surplus.

*Email: charlottegrace23@gmail.com. I thank my supervisors Paul Klemperer and Howard Smith for their advice and contributions, which have been critical at every stage of this project; it could not have been written without them. I am also grateful to Charlotte Barton, Alan Beggs, Ian Crawford, Rand Fakhoury, Julia Giese, Tom Horn, Dennis Jeevarajasingham, Jakub Kastl, Rohit Lamba and Rhys Phillips for helpful comments. The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees.

1 Introduction

This paper provides a method to test the relative performance of alternative models of bidders' behavior in auctions. It uses data on bids but does not require data on bidders' values. I analyse the Bank of England's (BoE) Indexed Long-Term Repo (ILTR) auctions, in which the central bank provides liquidity insurance to financial institutions by lending a quantity of funds against collateral. In these auctions, a bid expresses the price a bidder is willing to pay as a function of the quantity they demand. The payment rule is uniform-pricing. A bidder's bid for a marginal unit therefore may affect the price they pay for inframarginal units. This creates an incentive to bid differently from the bidder's true value for each unit of the good.

My main finding is that bidding behavior is nevertheless better explained by a model of "truthful bidding" than by a conventional model in which bidders are both strategic and risk neutral. In a truthful bidding model, bidders submit bid functions corresponding to their true marginal valuation functions. Conversely, in a conventional strategic model, bidders best respond to their beliefs about the economic environment and other bidders' behavior.

So the findings of this paper suggest that Milton Friedman (1960) may have been correct in advocating the uniform-price auction to issue U.S. Treasuries based on the view that bidders "need only know the maximum amount [they] are willing to pay for different quantities" (Friedman, 1991). He argued that bidders do not need to strategize so the uniform-price auction should reduce collusion and increase participation relative to the pay-as-bid auction.[1]

The difference between bidders' true values and their bids, i.e. their "bid shading", is difficult to measure because their values are rarely observed. Studies of divisible-good auctions therefore must assume a model of bidding when assessing the effi-

[1]Uniform-price and pay-as-bid are the two payment rules most commonly used in practice. In pay-as-bid auctions, a bidder pays their bid for each unit that they win. Their bid therefore also affects both their probability of winning and the expected payment.

ciency, revenue and other outcomes of alternative auction designs (see, e.g., surveys by Hortaçsu and McAdams, 2018; Hortaçsu and Perrigne, 2021). The nature of bidding behavior affects these market outcomes so the choice of model is critical to the analysis.

Truthful bidding can be justified as rational in two ways: risk aversion and costs of calculating the optimal strategy. First, when bidders are sufficiently risk averse, optimal strategies in the conventional model approximate truthful bidding, and I show that optimal strategies in the conventional model get closer to truthful bidding as risk aversion increases. (If bidding truthfully, bidders submit bids that correspond to their marginal values regardless of their degrees of risk aversion.) Second, truthful bidding is a best response when bidders face a cost of calculating what otherwise would be the optimal strategy and this cost outweighs the additional expected surplus generated by the strategy.

In this paper, I provide a test of competing models of bidding. It adapts Backus et al.'s (2021) test of firm conduct, which combines Rivers and Vuong's (2002) comparison of non-nested alternatives and Berry and Haile's (2014) exclusion restrictions. The procedure uses instrumental variables to evaluate the relative amount of bid shading that the different models predict. The instruments are variables that affect a bidder's bid shading but not their true value for liquidity. Under a correctly specified model, bidders' estimated values therefore will be uncorrelated with the amount of bid shading predicted by the instruments. Conversely, they will covary with the predicted bid shading under a misspecified model. Stronger covariance between the instruments and model-implied values is therefore evidence of worse model fit.

The key assumption for the test to be valid is that bidders have independent private values. Under this assumption, I use variables that measure the strength of competition, e.g. the number of other bidders in the auction, as instruments. One advantage of the approach is that, providing this assumption holds, the test remains valid even if both of the competing models are misspecified. This is because the procedure tests the models' *relative* performance. For example, if bidders are strategic but have different information to what is assumed in the conventional strategic model that I consider, the procedure will assess whether truthful bidding or the conventional model is a better approximation. While I focus on uniform-price divisible-good auctions, the method can be applied more broadly to auctions in which bidders' marginal values can be point identified and estimated under the candidate models of bidding behavior, and valid instruments are available.

I compare truthful bidding to Kastl's (2011, 2012) conventional share auction model, which modifies Wilson's (1979) foundational model to reflect the fact that bid functions are step functions and which is now widely used in empirical studies (e.g. Cassola et al., 2013; Hortaçsu et al., 2018; Allen et al., 2022). This is the most appropriate benchmark to analyse behavior in the BoE's liquidity auctions in which bidders submit a vector of ordered price-quantity pairs which constitute their stepped bid functions. In the auctions, bidders submit bid functions with very few steps, which is difficult to explain without Kastl's (2011) modification. Participants may optimally bid above or below their marginal values for liquidity in Kastl's (2011) model, so truthful bidding is an especially natural comparator.

I consider two explanations for why a model of truthful bidding outperforms this conventional model of behavior, which assumes bidders are both risk neutral and strategic: risk aversion and costs of determining the optimal strategy.

Kastl's (2011) conventional model of bidding assumes that bidders are risk neutral. However, there are two reasons why risk aversion might be a more appropriate assumption in the BoE's liquidity auctions. First, in a principal-agent framework, the manager tasked with bidding on behalf of the auction participant may have a concave utility function either because they are individually risk averse or because of the remuneration structure.[2] Second, the nature of the loans being allocated in the auctions suggests that financial institutions themselves might be risk averse—the auctions were introduced to respond to "the new demands for liquidity insurance that [the financial crisis] engendered" (Fisher, 2011a, p.15). Unsurprisingly, when bidders are risk averse, optimal strategies in the conventional model are closer to truthful bidding—if a bidder is sufficiently risk averse, they do not want to gamble on losing at prices at which they would prefer to win (or winning at prices at which they would prefer to lose) by submitting a bid that differs from their valuation. I build on Kastl (2011) by deriving novel identifying conditions, which allow for constant absolute risk aversion, in order to test the relative performance of the conventional model with varying degrees of risk aversion. I also show the best response gets closer to truthful bidding as risk aversion increases. The results show that truthful bidding can be rationalized within the conventional model by a degree of risk aversion which is consistent with that found in studies that are the most similar to my setting (Armantier and Sbaï, 2006; Boyarchenko et al., 2021).

Second, truthful bidding may be explained by a cost of determining the optimal strategy, i.e. a cost of "sophistication" (Hortaçsu and Puller, 2008), providing the cost exceeds the expected gains from what otherwise would be the optimal strategy. This resembles the finding that truthful bidding is an ε-equilibrium in Swinkels (2001) and Chakraborty and Englebrecht-Wiggans (2006): the loss from bidding truthfully rather than optimally becomes arbitrarily small as the number of bidders increases in multi-unit uniform-price auctions. I estimate a lower bound on the size of these costs by the difference between the surplus from what would be an approximate best response and the surplus from truthful bidding, assuming observed bids correspond to bidders' values. I find that the average lower bound is around 5% of bidder surplus, implying that, if bidders' actual strategies are to bid truthfully, they obtain up to 95%

[2]For example, this may result if the remuneration structure rewards not just higher expected profit for the participant but also winning at any price that the participant is willing to pay.

of the surplus that would have been generated by an approximate best response. In comparison, in Hortaçsu and Puller (2008), bidders' actual strategies only generate between 0% and 80% (excluding loss-making bidders) of the surplus generated by the optimal response.

This paper primarily contributes to the literature on testing equilibrium behavior in divisible-good auctions. Wolak (2007) and McAdams (2008) use violations of over-identifying moment conditions, which rule out profitable unilateral deviations in bidders' strategies. By contrast, I exploit exclusion restrictions to compare the goodness-of-fit of alternative models of behavior. I therefore do not require over-identifying moments, but I do require valid instruments. In this respect, the methods are complementary, allowing greater flexibility to the data available. By assessing relative rather than absolute model fit, my approach remains valid even if both models are misspecified. It has been used to test firm conduct in product markets (Backus et al., 2021; Duarte et al., 2023; Starc and Wollmann, 2022) and labor markets (Roussille and Scuderi, 2022). To my knowledge, it has not yet been adapted to an auction setting.[3]

Other studies estimate marginal values using data outside the model and compare these directly to the observed bids. For example, Wolfram (1999) and Borenstein et al. (2002) cannot reject strategic bidding in British and Californian electricity markets, whereas Hortaçsu and Puller (2008) and Hortaçsu et al. (2019) find strategic heterogeneity in Texas. In these later studies, larger bidders submit bids close to those predicted by profit maximization, yet smaller bidders persistently deviate from equilibrium bidding. Similar to my paper, Hortaçsu and Puller (2008) show that this can be rationalized by a cost of sophistication that only larger bidders are willing to incur.[4]

[3]Roussille and Scuderi (2022) model their labor market setting as an auction, in which firms submit wage offers, or "bids", for job candidates on an online job board. The structure of their data and their choice of instruments are similar to mine: individual firms' wage offers (i.e. bids) are observed and the number of candidates relative to the number of competitors (i.e. auction supply relative to demand), which measures firms' expectations about competing bids, is used as the instrument.

[4]Hortaçsu et al. (2019) demonstrate that a cognitive hierarchy model provides a good fit in the Texan electricity market, with strategic sophistication increasing in firm size. Truthful bidding is of course a

In contrast to these studies, my testing procedure does not require data on marginal values, which are often unavailable.

My paper also contributes to the literature on risk aversion in financial markets. Armantier and Sbaï (2006) estimate the degree of risk aversion of bidders in French Treasury auctions and Boyarchenko et al. (2021) calibrate the risk aversion parameter for bidders in US Treasury auctions. Applying the findings of these studies to the BoE's liquidity auctions would suggest that truthful bidding is approximately optimal in our context. In contrast, Allen and Wittwer (2023) estimate a much smaller degree of risk aversion when focusing exclusively on primary dealers in Canadian Treasury auctions.[5] Using data from Indian Treasury bill auctions, Gupta and Lamba (2023) use risk aversion to explain the combination of increased participation and higher uncertainty during the 2013 "taper tantrum". The main differences between my model and theirs are that I assume bidders have constant absolute risk aversion (CARA) rather than constant relative risk aversion (CRRA), and I allow for bidders to tie (i.e. for bidders to submit the same bid price). Ties are highly probable in my context and in others in which the set of possible bid prices is discrete.[6]

The paper proceeds as follows. Section 2 provides background for the BoE's liquidity auctions and introduces the bidding data. Section 3 describes the alternative models of behavior. Section 4 explains the identification of bidders' values, the estimation method and the testing procedure. Section 5 presents the results and shows how truthful bidding can be explained. Section 6 concludes.

special case in which Level-0 behavior is defined by truthful bidding and adopted by all bidders.

[5] Outside of financial markets, Häfner's (2023) analysis of meat quota auctions is the only study, to my knowledge, of risk aversion in multi-unit auctions.

[6] CARA implies that a bidder's risk aversion does not vary with their net surplus from the auction (defined by the area between their marginal valuation function and the auction price, up to the quantity they are allocated), whereas Gupta and Lamba's (2023) CRRA specification implies that the bidder's risk aversion is decreasing in their net surplus from the auction. CARA seems more appropriate in my context for two reasons. First, the financial institution's net surplus from the auction is trivial relative to its total assets so CARA seems a better approximation if the financial institution itself is risk averse. Second, if the manager tasked with bidding on behalf of the participant is individually risk averse, CRRA would imply that their risk aversion is lower when the stakes are higher.

2 Institutional Setting and Data

2.1 Indexed Long-Term Repo Auction

My setting is the Bank of England's (BoE) Indexed Long-Term Repo (ILTR) auction. It was introduced in 2010 to efficiently respond to "the new demands for liquidity insurance that [the financial crisis] engendered" (Fisher, 2011a, p.15) by lending funds to financial institutions against multiple types of collateral. It allocates the largest amount of funds among the BoE facilities that provide liquidity insurance. I study the period June 2010 – May 2012, in which the auctions were held monthly.[7,8]

Loans of central bank reserves are issued at a spread over the BoE base rate ("Bank Rate") for a 3- or 6-month term, and separate operations are run for different terms. The lending is collateralized. Within each operation, bidders may borrow the reserves against one of two types of collateral: "Level A", including highly liquid gilts, sterling Treasury bills and certain sovereign and central bank debt, and "Level B", including high quality, but less liquid, sovereign debt and certain asset-backed securities.[9] The ILTR operations therefore provide a "liquidity upgrade" to participants, allowing them to swap collateral for more liquid reserves. I refer to funds borrowed against each type of collateral by goods A and B respectively, and the spreads over Bank Rate at which the loans are allocated by the prices of goods A and B.

Each operation is a Product-Mix Auction, originally developed by Klemperer (2008)

[7] I limit the sample to the first 24 months of the ILTR auctions because total aggregate demand of the participants decreases significantly after this period (see Grace (2024) for further discussion). This can be explained by an increase in the availability of funding from the BoE's quantitative easing programme, a reduction in participants' holdings of assets in the less liquid collateral set, and liquidity regulation which required funding of longer maturity than the 3-month term auctions (Winters, 2012). Increased availability of funds from the euro area and alternative BoE facilities also likely played a role.

[8] The auction rules, supply curves representing the BoE's preferences between collateral sets, and eligibility for entry have evolved since. Further details are at https://www.bankofengland.co.uk/markets/bank-of-england-market-operations-guide/our-tools.

[9] The classification of collateral sets reflected the relative liquidity of the assets (Fisher, 2011b), and the BoE imposed haircuts on the assets with the aim that differences between the collateral sets reflect only these liquidity premia, not credit risk premia. The haircuts applied to assets varied depending on, e.g. credit rating, interest rate, and maturity (Bank of England, 2010a). The classification was adjusted over the period (see Bank of England, 2011). The current classification can be found at https://www.bankofengland.co.uk/-/media/boe/files/markets/eligible-collateral/summary-table-of-collateral.pdf.

and further described in Klemperer (2010, 2018). It is a sealed-bid uniform-price auction in which goods are jointly allocated, as explained below. For each good, all winning bidders pay the auction price for that good.

Participants of the BoE's Sterling Monetary Framework (SMF) with access to the BoE's Open Market Operations (OMO) were eligible to register to participate in the ILTR auctions, including banks and building societies. The number of institutions with OMO access ranged from 48 to 52 in the sample period. There was no obligation for registered bidders to participate in the auctions.

Bidders can submit any number of bids. A bid specifies the good (i.e. the type of collateral used), the quantity demanded, and the price (i.e. spread) the bidder is willing to pay. For example, a bid may specify demand of £50 million at a spread of 2 basis points (bps) for good A.[10] The minimum bid price is 0bps, with increments of 1bp. The minimum bid size is £5 million, with increments of £1 million, and the minimum unit of allocation is £100,000. A bidder may bid for a maximum of £1.5 billion and £0.75 billion of loans for 3-month and 6-month terms, respectively. For each good, a bidder's bids imply a (price, quantity)-schedule, i.e. an individual bid function, which is a step function.

Prior to the auction, the BoE commits to (i) a publicly announced maximum supply made available across goods A and B, and (ii) a privately known "relative supply" curve. The maximum supply was £5 billion for the 3-month term auction and £2.5 billion for the 6-month term auctions. The relative supply is "pinned down by [the BoE's] preferences" to provide liquidity insurance at prices which incentivize prudent liquidity management (Fisher, 2011a, p.12). It is measured by the spread between the prices of goods B and A as a function of the quantity allocated of good B. The minimum spread was 5bps, except for in the 6-month auctions from May 2011, in which

[10] Bidders were also permitted to submit "paired bids", which specify the quantity demanded and a price for each of goods A and B. For each bid, the bidder may be allocated the loan against either good A or good B. Because bidders rarely made use of paired bids, I drop them from the sample.

Table 1: Summary statistics across both goods

	3-month term		6-month term	
	Mean	SD	Mean	SD
Maximum supply (£billion)	5	0	2.5	0
Total quantity of funds demanded (£billion)	4.57	2.67	2.69	1.93
Number of bidders	8.13	3.07	7.63	3.85
Total quantity of funds allocated (£billion)	3.55	1.86	1.65	0.99

the minimum spread was 15bps. The relative supply curves were flat at the minimum spreads up to some fixed quantities of good B and beyond this were increasing; I am not permitted to disclose their precise functional forms.

The Product-Mix Auction uses the information from the submitted bids and auctioneer's supply curves to find the competitive equilibrium, assuming bids correspond to bidders' marginal values and the supply curves represent the prices the BoE is willing to accept. The quantity allocated of each good therefore depends on the submitted bids and BoE's supply preferences, and so is uncertain from the perspective of bidders (see Appendix A of Grace (2024) for an illustration).

For each good, the auction price is the maximum of the highest losing bid for that good and the minimum price the BoE is willing to accept, as expressed by the supply curves.[11] Bids strictly above the auction price are fully allocated and bids on the margin are rationed pro-rata for each good. Winning bidders pay the auction prices.

2.2 Bidding Data

I use a unique proprietary dataset of all bids submitted in the 16 auctions of 3-month term and the 8 auctions of 6-month term in June 2010 – May 2012, summarized in Table 1. To my knowledge, this data has only once been analysed.[12]

Table 2 provides summary statistics by good. The auction price of good A is typ-

[11]For good B, the minimum price the BoE is willing to accept is the sum of the auction price for good A and the relative supply curve evaluated at the quantity allocated of good B. For good A, it is the difference between the auction price for good B and the relative supply curve evaluated at the quantity allocated of good B.

[12]de Roure and McLaren (2021) study the risk profile of participants in the ILTR auctions.

Table 2: Summary statistics by good

	Good A				Good B			
	3-month term		6-month term		3-month term		6-month term	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Quantity of funds demanded (% of allocation)	125.72	22.68	171.98	66.62	151.26	58.15	202.74	91.49
Bid price (weighted, basis points)	2.51	2.85	3.06	4.31	23.61	9.96	38.96	16.15
Number of bids per bidder (unweighted)	1.52	0.84	1.49	0.86	1.64	1.31	2.30	1.71
Number of bids per bidder (weighted)	1.79	0.94	1.86	1.04	2.90	2.11	2.95	2.28
Quantity demanded by bidder (% of maximum supply)	13.65	11.87	13.93	12.30	4.30	6.41	9.69	9.87
Quantity of funds allocated (% of total allocated)	83.72	12.00	70.65	20.19	16.28	12.00	29.35	20.19
Price paid (weighted, basis points)	1.14	1.71	0.59	0.79	21.67	6.52	43.43	1.37

Note: The quantity of funds demanded for each good is expressed as a percentage of the quantity allocated of that good; quantities of funds demanded and quantities of funds allocated per good are weighted by the aggregate quantities demanded in auctions; prices paid for each good are weighted by the aggregate quantities allocated of the good in auctions; bid prices are weighted by the quantities demanded at those prices by bidders; bps denote basis points.

ically close to its reserve of 0bps (the weighted average price paid is 1.14bps in the 3-month term auctions and 0.59bps in the 6-month term auctions), reflecting the high liquidity of assets in the Level A collateral set. In contrast, the assets in the Level B set are more heterogeneous and typically have limited secondary markets. This is reflected in much higher weighted average prices paid for good B of 21.67bps and 43.43bps in the 3-month and 6-month term auctions, respectively, and much greater variance in the auction prices than good A.

2.3 Bidding Behavior by Bidder Size

In conventional models of bidding in uniform-price auctions, a bidder can reduce the expected auction price paid for all units that they win by reducing the quantity that they demand at each price; the strength of this incentive is increasing in the quantity that the bidder demands. This corresponds to the quantity-shading incentives of oligopolists facing uncertain demand in Klemperer and Meyer (1989). If bidders were acting strategically, we would therefore expect to observe larger bidders to *ceteris paribus* submit lower bid prices relative to their values.

Table 3 shows the average bid prices of "Small" bidders (those who demand less than 2.5% of the maximum supply in all auctions) and "Large" bidders (the remaining set). (Natural alternative definitions of bidder size give similar results.) Larger bidders

Table 3: Summary statistics by bidder size

	Good A								Good B							
	3-month term				6-month term				3-month term				6-month term			
Bidder size	Small		Large		Small		Large		Small		Large		Small		Large	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Bid price (bps)	3.22	2.13	2.51	2.85	5.33	5.20	3.03	4.29	24.42	13.93	23.47	9.10	41.36	23.49	38.88	15.85
# of bids per bidder	1.09	0.29	1.67	0.91	1.25	0.45	1.58	0.96	1.28	0.72	2.14	1.74	1.89	1.17	2.50	1.92

Note: Small bidders are defined as those who demand less than 2.5% of the maximum supply in all auctions in the sample period; Large bidders are defined as those in the remaining set; bid prices are weighted by the quantities demanded at those prices by bidders, and numbers of submitted bids are unweighted; bps denote basis points.

do submit lower bid prices in the ILTR auctions, but this can be explained by variation in bidders' strategic incentives or variation in their values. It is possible, for example, that larger bidders in the ILTR auctions have greater access to alternative funding sources which affect their valuations. The main contribution of this paper is to employ a testing procedure, described in Section 4.3, which isolates the variation in bidders' strategic incentives in order to evaluate alternative models of behavior.

2.4 Bids Submitted per Bidder

A striking feature of the data is that bidders submit few bids per auction (Figure 1). This is also seen in many other multi-unit auction settings.[13] It motivates the strategic model, developed by Kastl (2011, 2012), that I use as a benchmark. In the model, bidders face costs of submitting additional bids, so the number of bids to submit is a strategic choice: a bidder must trade off the marginal "bidding cost" with the marginal benefit of fine-tuning their demand by submitting an additional bid. If the marginal benefits of fine-tuning are small, bidders optimally submit few bids.

Larger bidders submit more bids per good (see Table 3), but the number of bids they submit is still small. Without bidding costs or other frictions, this is difficult to explain within a strategic framework. Even if a bidder had a constant marginal value for the

[13] In Czech, Canadian and US Treasury auctions, bidders on average submit 2.3 bids (Kastl, 2011), 2.9 bids (Hortaçsu and Kastl, 2012) and between 3 and 5 bids (Hortaçsu et al., 2018), respectively. In central bank operations, Cassola et al. (2013) find that participants submit 1.7 bids on average in the European Central Bank's auctions for short-term repurchase agreements, and 82% of participants in the Bank of Canada's cash management auctions submit only one bid (Chapman et al., 2006).

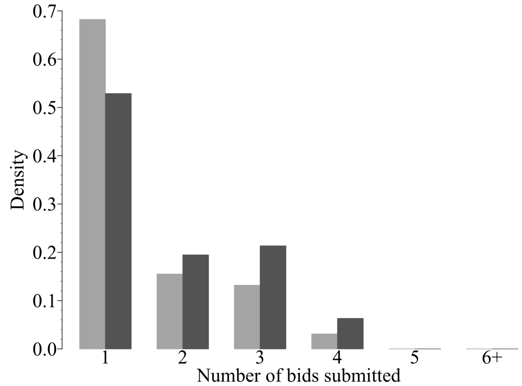

(a) Good A

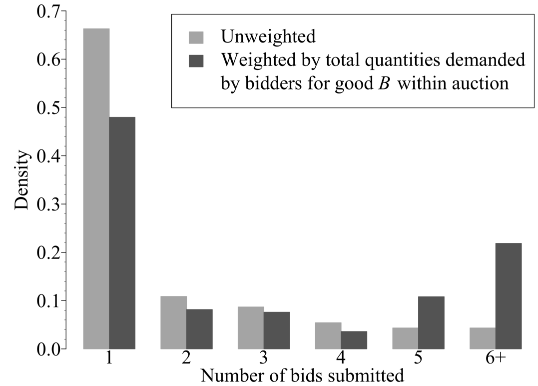

(b) Good B

Figure 1: Number of bids submitted per bidder within auction

good, we would expect them to submit a large number of bids, as their incentive to understate their value is increasing in the quantity that they demand (see Section 2.3).

3 Candidate Models of Behavior

This paper analyses competing models of bidders' behavior using a testing procedure based on instrumental variables. In this section, I first lay out the assumptions that are required for the test to be valid; the main assumption is that bidders have independent private values. I then describe the competing models.

The first model is a "truthful bidding" model in which all bidders submit bids according to their true marginal values. The second is Kastl's (2011) conventional strategic model, which captures the key features of the BoE's liquidity auctions, including that each bidder makes few bids per auction. One limitation is that the model does not permit bidders to be risk averse, so I relax Kastl's (2011) assumption of risk neutrality. This generalisation shows that optimal strategies in the strategic model tend towards truthful bidding as the degree of risk aversion becomes large. The correct model may differ from both of these models; the procedure tests which of the two is closer.

Section 3.4 describes the amount of "bid shading"—the difference between bidders' marginal values and their bids—implied by the different models. In the truthful bidding model, bid shading is of course always zero, whereas in the conventional model it varies across bidders and across the quantities they demand. This variation is key to the test's identification. The following description refers to one auction and I suppress the time subscript.

3.1 Assumptions on the Economic Environment

Two goods, A and B, are perfectly divisible, measured in the same units, and allocated by a Product-Mix Auction as described in Section 2.1.

There are N potential bidders. Prior to the auction, each bidder $i \in \mathcal{N} = \{1, ..., N\}$, receives a privately known, multi-dimensional signal, Θ_i. I assume each bidder's signal is distributed such that their value is above the reserve price for at most one of the goods, so they only bid on at most one good. This greatly improves the tractability of the model and is appropriate for the majority of bidders (85%) in the data who only bid for one good. For the purpose of the analysis, it is also appropriate for an additional 4% of bidders per auction who bid on both goods, but for at least one of the goods, they submit only one bid and that bid is at the reserve price.[14,15] Let $\boldsymbol{\Theta_{-i}} = \{\Theta_j\}_{\{j \neq i\}}$ denote the signals received by all bidders except i.

Motivated by the evidence of heterogeneity described in Section 2.3, bidders for each good are split into two groups, so the set of potential bidders for good $g \in \{A, B\}$ with bidder size $h \in \{1, 2\}$ is $\mathcal{N}^{g,h}$. This set has dimension $N^{g,h}$, where $\sum_g \sum_h N^{g,h} = N$. Bidders are assumed to be symmetric within group, with marginal

[14] The analysis uses local deviations in the quantity that a bidder demands to identify their values. If a bidder bids on both goods, but their bid is at the reserve for good A, then the bidder is always indifferent between winning and losing good A, so a deviation in the quantity they demand of good B does not affect their surplus from good A (and vice versa). We can therefore treat a bidder for both goods as if they are two bidders, one for each good.

[15] Moreover, for those bidders who do submit bids for both goods within an auction, additive separability of their utility functions in the quantities allocated of the two goods seems a reasonable assumption. This is because bidders have the option to submit "paired bids", which allow bidders to express rich one-for-one substitutes preferences between goods A and B (see Klemperer, 2018), but they very rarely choose to submit such bids. Complementarity seems implausible (the BoE designed the ILTR auction anticipating bidders to express one-for-one substitutability between goods A and B, if at all).

value functions that are non-increasing and continuous in the allocated quantities.[16]

The main assumption required for the testing procedure to be valid is that bidders have independent private values. Let $y_i \in \mathbb{R}_+$ be the quantity allocated to bidder i.

The IPV Assumption *The marginal valuation of bidder* $i \in \mathcal{N}^{g,h}$, $g \in \{A, B\}, h \in \{1, 2\}$, *denoted* $v^{g,h}(y_i, \Theta_i)$, *is only a function of* y_i *and the bidder's signal* Θ_i.

Three reasons suggest that an independent private values framework is a reasonable approximation. First, while the ILTR auctions were designed to meet demands for liquidity insurance, the sample period of 2010 – 2012 was relatively stable. This suggests that the interdependence that would arise from another bidder's private signal revealing information of perceived market stress was less relevant in our sample.

Second, it is plausible that participation over the period studied was driven largely by idiosyncratic liquidity needs. The relatively short maturity of loans secured in the ILTR auctions limited incentives for speculation,[17] especially for good B, for which collateral assets were less liquid by definition. Moreover, the general collateral repo rate was a good proxy for the market value of good A. Because this rate was readily observable, the uncertainty about the future market value of good A was plausibly symmetric, which is consistent with an independent private values framework.

Finally, neither Bindseil et al. (2009) nor Hortaçsu and Kastl (2012) find evidence for interdependent values in ECB repo auctions and Canadian Treasury auctions, respectively. It is not possible to perform these tests in our context, but their findings are informative given the similarity of our context to theirs.

Turning to the utility function, bidders are assumed to have utility functions which

[16]The funds that a given bidder is allocated in the auction is one component of their larger stock of cash reserves. The setup costs are mostly sunk as collateral for ILTR loans is typically deposited at the BoE prior to the auction, and there are no other obvious fixed costs, which suggests limited scope for increasing returns. If any allocated funds are put to its highest value use, a non-increasing marginal valuation function is a reasonable assumption.

[17]Hortaçsu and McAdams (2010) and Haile (2001) provide evidence, in the Turkish Treasury security market and timber market respectively, of limited scope for resale (i.e. speculation on the future price) at short time horizons.

are quasi-linear in assets outside the auction. I allow for bidders to be risk averse, with constant absolute risk aversion utility described by parameter $\rho^g \geq 0$, which is common to all bidders for good $g \in \{A, B\}$. Let P^g be the auction price of good $g \in \{A, B\}$. The expected utility of type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ given strategy profile $\boldsymbol{\sigma}(\boldsymbol{\Theta}) = \{\sigma_j(\Theta_j)\}_{\{j \in \mathcal{N}\}}$ in Model m is therefore

$$V_i(\theta_i) = \begin{cases} \mathbb{E}\Big[U_i\Big(y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})), P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))|\theta_i\Big)\Big] & \text{if } \rho^g = 0 \\ \mathbb{E}\left[\frac{1}{\rho^g}\Big(1 - e^{-\rho^g U_i(y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})), P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))|\theta_i)}\Big)\right] & \text{if } \rho^g > 0 \end{cases}$$

where $U_i\big(y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})), P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))|\theta_i\big) = \int_0^{y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)} v^{g,h}(u, \theta_i)du - P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i) y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)$.

Bidder i submits K_i bids, which are summed in (price, quantity)-space to form an individual bid function. For each bid $k \in \{1, ..., K_i\}$, the bid price is $b_{i,k}$, and the cumulative quantity demanded is $q_{i,k}$, where the bids are ordered in increasing quantity. Let $\boldsymbol{b_i} \in \mathbb{R}_+^{K_i}$ and $\boldsymbol{q_i} \in \mathbb{R}_+^{K_i}$ respectively denote the vector of bid prices and cumulative quantities demanded by bidder i. I refer to $(b_{i,k}, q_{i,k})$ as the bid price and quantity demanded at step k in bidder i's bid function. Corresponding to the ILTR auction rules, bids must be whole basis points, and I approximate the permitted grid of quantities by assuming that quantities demanded are continuous in the unit interval.

3.2 The Truthful Bidding Model

The first of the two competing models is a "truthful bidding" model, in which all bidders optimally submit bids which approximate their true demand for liquidity.[18]

Definition 1 (Model *T*) *In Model T, the IPV Assumption holds and each bidder's bid function corresponds to the greatest integer function of their marginal valuation function (i.e. for all quantities, the largest integer weakly below their marginal value).*

[18] This can be motivated in many ways, e.g. each bidder faces an arbitrarily large cost of determining what would otherwise be the optimal bid function. Truthful bidding is a natural heuristic in the strategic model that I consider because the optimal bid may be above or below a bidder's true value.

3.3 The Conventional Strategic Model

The second model closely follows Kastl (2011), except for I relax the assumption of risk neutrality. In this model, bidders choose their bids to maximize their expected surpluses, conditional on the information available to them at the time of bidding.

I assume bidders have common knowledge of the number of potential bidders, the joint distribution of signals, and the risk aversion parameters, whereas the number of other bidders and their bids are unknown.[19] In the ILTR auctions, the suitability of this common prior is unclear. Some information about the number of potential bidders was publicly available during the period (e.g. Winters, 2012). Bidders also may have learned about the distribution of other bidders' values in auctions held by the BoE prior to our sample period, which had similar purpose to the ILTR auctions but differed in design.[20] In addition, for simplicity, I assume that the BoE's maximum supply and relative supply curve are commonly known.[21]

Different assumptions on bidders' information would imply different optimal strategies. One of the benefits of the testing procedure is that it tests the *relative* performance of the two models and so is robust to their misspecification. For example, if bidders are strategic but have different information from what is assumed here, the procedure tests which of the two models are a better approximation of this correct model.

In the model, bidders face costs of submitting additional bids, which are increasing in the number of bids that they submit. Kastl (2011) introduces these "bidding costs", denoted $c_i(K_i)$, to rationalize the fact that bidders submit few bids per auction (see Section 2.4). The costs can be interpreted as the physical and time costs of submitting

[19] Common knowledge of the number of potential bidders can be relaxed by instead assuming that the number of potential bidders is stochastic, but its distribution is common knowledge. Hortaçsu (2002) finds that this modification does not have a significant impact.

[20] The ILTR auctions replaced the BoE's Long-Term Repo operations, which were pay-as-bid auctions in which a reserve price of 50bps was imposed for bids on good B (see Fisher, 2011a).

[21] We could relax this assumption by assuming that the distribution of supply of each good is commonly known or we could assume that the BoE is a strategic actor with privately known preferences. Hortaçsu (2002) finds this makes very little difference to his results.

bids or of fine-tuning them (see Grace (2024) for discussion).

Definition 2 (Model $S(\rho)$) *In Model $S(\rho)$, the IPV Assumption holds; the number of potential bidders in each group ($N^{g,h}, g = \{A, B\}, h = \{1, 2\}$), the joint distribution of signals, the risk aversion parameters ($\rho = (\rho^A, \rho^B)$), the maximum supply and the relative supply curve are commonly known; and $c_i(K_i)$ is the cost of submitting bids, where $c_i(K_i + 1) \geq c_i(K_i) \geq 0 \, \forall K_i \in [1, \bar{K} - 1]$.*

The solution concept is a group-symmetric Bayesian Nash Equilibrium. An equilibrium is a strategy profile, $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, such that for every bidder $i \in \mathcal{N}$ and almost every type θ_i, $\sigma_i(\theta_i)$ solves $\sigma_i(\theta_i) \in \arg\max_{\sigma_i(\theta_i) \in \mathcal{F}_i} \left(V_i(\theta_i) - c_i(K_i)\right)$

3.4 Bid Shading in Equilibrium

The testing procedure described in Section 4.3 exploits differences in the amount of "bid shading"—the difference between bidders' marginal valuations and their bids—implied by the alternative models of behavior in order to compare their model fit.

Definition 3 (Bid shading) *The bid shading of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ at the quantity they demand at step $k \in \{1, ..., K_i\}$ is $\mu_{i,k} = v^{g,h}(q_{i,k}) - b_{i,k}$.*

In Model T, an equilibrium strategy profile is one in which each bidder submits a set of bids which correspond to the greatest integer function of their marginal valuation function, regardless of their degree of risk aversion. Because the marginal valuation function is assumed continuous, this implies that a bidder's bid at step $k \in \{1, ..., K_i\}$ is equal to their marginal value so that $\mu_{i,k} = 0 \, \forall k \in \{1, ..., K_i\}, i \in \mathcal{N}$.

Bid shading in Model $S(\rho)$ is more complicated. A necessary condition for a strategy profile to be a Bayesian Nash Equilibrium in this model is that bidder $i \in \mathcal{N}$ cannot increase their expected utility by unilaterally deviating from their prescribed strategy, σ_i, given their beliefs. Each bidder's strategy is multi-dimensional so there are many possible ways in which to deviate.

The main identifying condition for bid shading rules out profitable deviations in the quantity that a bidder demands at a particular step for type $\theta_i \in \Theta_i$ in Model $S(\rho)$, holding the rest of their strategy constant. Proposition 1 (below) states this condition, generalising Kastl's (2011) corresponding condition to allow for risk aversion.

The benefit of this deviation to the bidder is the difference between two effects. The first effect (the left-hand side of Equation 1) is the increase in expected utility from winning the marginal unit at step k, holding the distribution of the auction price constant. Following the standard Kastl (2011) decomposition, this distinguishes between the cases in which the bidder wins the full marginal unit, and the bidder "ties" with other bidders at step k and step $k+1$ and so is rationed.

The second effect (the right-hand side of Equation 1) is a market power effect. It measures the reduction in expected utility due to the bidder having to pay more for the units that they win if the increase in their demand increases the auction price.

A necessary condition for equilibrium in Model $S(\rho)$ is that the bidder does not benefit from deviating, i.e. the two effects must be equal.

Proposition 1 (Necessary condition on quantity deviations in Model $S(\rho)$) *In Model $S(\rho)$ with $\rho = (\rho^A, \rho^B) \geq 0$, in any type-symmetric Bayesian Nash Equilibrium, for almost every type θ_i, every step k in the K_i-step bid function in strategy σ_i of bidder $i \in \mathcal{N}^{g,h}$, $g \in \{A, B\}, h \in \{1, 2\}$, must satisfy*

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho^g U_i}(v^{g,h}(q_{i,k}, \theta_i) - P^g)\middle| P^g \in (b_{i,k+1}, b_{i,k})\right] \mathbb{P}(P^g \in (b_{i,k+1}, b_{i,k})) \\
&+ \sum_{x = b_{i,k}, b_{i,k+1}} \left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho^g U_i}\left(v^{g,h}(y_i, \theta_i) - x\right) \tfrac{\partial y_i}{\partial q_{i,k}} \middle| P^g = x \wedge Tie^g\right] \mathbb{P}(P^g = x \wedge Tie^g) \right) \\
&= \begin{cases} q_{i,k} \frac{\partial}{\partial \epsilon} \left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}} \left[\tilde{P}^g(\epsilon) \mathbb{I}\big(\tilde{P}^g(\epsilon) \in [b_{i,k+1}, b_{i,k}]\big) \right] \right) \Big|_{\epsilon=0} & \textit{if } \rho^g = 0 \\ \frac{1}{\rho^g} \frac{\partial}{\partial \epsilon} \left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}} \left[e^{-\rho^g U_i} \mathbb{I}\big(\tilde{P}^g(\epsilon) \in [b_{i,k+1}, b_{i,k}]\big) \right] \right) \Big|_{\epsilon=0} & \textit{if } \rho^g > 0 \end{cases} \qquad (1)
\end{aligned}$$

where $P^g = P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)$, $y_i = y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)$, and $U_i = U_i(y_i, P^g|\theta_i)$; $(P^g = p \wedge Tie^g)$

denotes the event that the auction price for good g is p and at least one other bidder submits a bid for good g with a bid price equal to p; $b_{i,K_i+1} = 0$; $\mathbb{I}(.)$ is the indicator function; and $\tilde{P}^g(\epsilon)$ is the auction price of good g if bidder i unilaterally deviates to a strategy in which type θ_i of bidder i demands a quantity $(q_{i,k} - \epsilon)$ for good g at step k and the rest of the strategy profile, including the rest of their strategy, is unchanged.

Proof. See Kastl (2011) for $\rho^g = 0$ and Online Appendix A.1 for $\rho^g > 0$. ■

To understand this condition, it is helpful to first consider the simplest case in which a bidder takes the distribution of the auction price as given and is risk neutral. In this case (ignoring ties for convenience), Equation 1 simplifies to $v^{g,h}(q_{i,k}) = \mathbb{E}\big[P^g \big| P^g \in (b_{i,k+1}, b_{i,k})\big]$; the bidder demands a quantity at step k which equates their marginal value for the good with the expected auction price, conditional on winning. Because the auction price is uncertain, the bidder's demand for the marginal unit is equivalent to accepting a lottery with an expected payoff of zero.

If the price-taking bidder is instead risk averse (and may tie), winning the marginal unit at a loss (i.e. at a price above their marginal value) has a larger, negative, impact on the bidder's utility than winning the unit at a gain. For the same reason that risk aversion makes bidders more aggressive in single-unit first-price auctions (Krishna, 2009), risk aversion reduces the difference between the bid price that a bidder submits and their marginal value. With sufficiently high risk aversion, any lottery with the possibility of a negative payoff is undesirable, so a price-taking bidder who faces ties on the margin will bid truthfully.[22]

More generally a bidder has market power, as the distribution of the auction price depends on the bids that they submit. Demanding a larger quantity weakly increases the price, which is paid for all units that they win.

[22]This result depends on the assumption of independent private values. In their study of Indian Treasury bill auctions, Gupta and Lamba (2023) assume that bidders' values have a common component. In their model, risk averse bidders increase their bid shading as uncertainty in the common value of the good rises to protect against the risk of a low ex post common value.

If the bidder is risk neutral, market power creates an incentive for bid shading that corresponds to the quantity-shading incentive of an oligopolist facing uncertain demand in Klemperer and Meyer (1989). In my setting, the bidder has an incentive to bid for a quantity below the true amount that they demand at each price, in order to lower the auction price and reduce the total amount they must pay. The incentive to do so is increasing in the quantity that they demand.

At moderate levels, risk aversion's effect on the incentive to exert market power depends on the distribution of the auction price. Risk aversion weakens the incentive if the bidder's market power is weaker at higher prices. This is because exerting less market power would reduce the dispersion of the auction price, and hence of the bidder's utility. Because a risk averse bidder prefers a less dispersed distribution of the auction price, risk aversion therefore *ceteris paribus* reduces the amount of bid shading of a bidder with market power in this case. The converse holds if the bidder's market power is weaker at higher prices.

Nonetheless, as risk aversion increases, the difference between the bidder's bid and marginal value decreases for the same reasons as for the price-taking bidder. At a sufficiently high degree of risk aversion, a bidder bids truthfully, regardless of their market power.

4 Econometric Method

This section explains how the models are compared. I first describe the identification (Section 4.1) and estimation (Section 4.2) of bidders' marginal values, and therefore their bid shading, in each model. Section 4.3 then describes the testing procedure which is based on the relative amount of bid shading implied by the models.

4.1 Identification

In Model T, the marginal value of bidder $i \in \mathcal{N}$ at the quantity that they demand at step $k \in \{1, ..., K_i\}$ of their bid function is equal to their bid at step k, $b_{i,k}$.

In Model $S(\rho)$, bidders' marginal values must satisfy Equation 1 in any type-symmetric Bayesian Nash Equilibrium. This condition can be used to recover bidders' unobserved marginal values at the quantities they demand at the steps of their bid functions from the observed bids, identifying their bid shading at these quantities.[23]

If bidder i does not expect to submit a bid at the same price as another bidder and therefore "tie" if their bid is marginal, their marginal value at the quantity they demand at step k can be readily estimated (see Equation 3 in Online Appendix A.2). However, ties are important in our setting because bidders frequently submit bids at the same prices as other bidders. For goods A and B, respectively, 69% and 20% of bids are submitted at prices at which at least one other bidder bids in the same auction.

Allowing for ties makes identification more complex because the bidder is rationed pro rata if the auction price equals their bid. To evaluate the benefit of the quantity deviation, we therefore need to know the quantity the bidder will win if they do tie, which depends on the rationing coefficient, and their marginal value function between the quantity they win and the quantity they demand.

I therefore make two assumptions in order to identify bidders' marginal values in Model $S(\rho)$. The first assumption is that bidder i's marginal valuation function is flat around the quantity that they demand, so that their marginal value at the quantity that they win when rationed is the same as their marginal value at the quantity they demand. For bidders who submit one bid (67% of bidders per auction in our sample), this corresponds to the assumption made by Kastl (2011) for identification when allowing for ties in Model $S(0)$.[24] Some natural alternative assumptions would achieve the same result, e.g. that the marginal value function is piecewise linear.

[23] Grace (2024) provides additional necessary conditions for equilibrium in Model $S(0)$ which can be used to set-identify bidders' full marginal valuation functions.

[24] Kastl (2011) assumes that the expected marginal valuation function in the range $y_i \in [\mathbb{E}[y_i|P^g = b_{i,k} \wedge Tie^g], q_{i,k}]$ is constant at $v^{g,h}(q_{i,k}, \theta_i)$. For bidders who submit more than one bid per auction, allowing for risk aversion also requires the stronger assumption that the marginal valuation function is also flat between the quantity that the bidder demands at step k and the quantity they are allocated when rationed at step $k+1$.

The second assumption is that the rationing coefficient that bidder i expects to face at price $b_{i,k}$, conditional on the strategy profile in Model $S(\rho)$, is deterministic. This approximates a situation in which the bidder does not take the variation in their net utility from the auction into account when evaluating the marginal benefit of winning an additional unit at a auction price equal to their bid.[25]

Under these assumptions, the necessary condition for bidders' strategies to be an equilibrium in Model $S(\rho)$ simplifies, as shown in Online Appendix A.3, to:

$$\begin{aligned}
&\mathbb{E}\Big[e^{\rho^g P^g q_{i,k}}\big(v^{g,h}(q_{i,k},\theta_i)-P^g\big)\Big|P^g\in(b_{i,k+1},b_{i,k})\Big]\mathbb{P}(P^g\in(b_{i,k+1},b_{i,k})) \\
&+e^{\rho^g((1-r^g_{i,k})(v^{g,h}(q_{i,k},\theta_i)-b_{i,k})+b_{i,k})(q_{i,k}-q_{i,k-1})}\tilde{v}_k\mathbb{P}(P^g=b_{i,k}\wedge Tie^g) \\
&+e^{\rho^g(r^g_{i,k+1}(q_{i,k+1}-q_{i,k})(b_{i,k+1}-v^{g,h}(q_{i,k},\theta_i))+b_{i,k+1}q_{i,k})}\tilde{v}_{k+1}\mathbb{P}(P^g=b_{i,k+1}\wedge Tie^g) \\
&=\begin{cases} q_{i,k}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(\tilde{P}^g(\epsilon)\in[b_{i,k+1},b_{i,k}]\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g=0 \\ \frac{1}{\rho^g}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}\left[e^{\rho^g\tilde{P}^g(\epsilon)q_{i,k}}\mathbb{I}\left(\tilde{P}^g(\epsilon)\in[b_{i,k+1},b_{i,k}]\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g>0 \end{cases}
\end{aligned} \tag{2}$$

where $\tilde{v}_k=\big(v^{g,h}(q_{i,k},\theta_i)-b_{i,k}\big)r^g_{i,k}$ and $\tilde{v}_{k+1}=\big(v^{g,h}(q_{i,k+1},\theta_i)-b_{i,k+1}\big)\big(1-r^g_{i,k+1}\big)$.

This can be solved numerically to recover bidders' marginal values, by choosing a risk aversion parameter and using the bidding data to estimate bidders' beliefs over the distribution of the auction price and the impact of their own bids on this distribution.

Because it is not possible to identify the degree of risk aversion from Equation 2, I test the relative performance of models which differ in this parameter in Section 4.3.

4.2 Estimation

I estimate each bidder's marginal values at the quantities they demand at the steps of their bid function separately for each auction in which they participate.

For Model T, the marginal value of bidder $i\in\mathcal{N}$ at the quantity they demand at

[25] We could relax the assumption by estimating the bidder's net utility for each realisation of the rationing coefficient and weighting by the probability of its occurrence. This would be difficult to estimate precisely given the available data.

step k of their bid function is trivially estimated by their bid at step k, $b_{i,k}$.

For Model $S(\rho)$, we must estimate bidders' beliefs over the distribution of the auction price in order to estimate their values. Under the assumption that bidders are playing a Bayesian Nash Equilibrium, each bidder forms beliefs over the distribution of auction prices that are consistent with other players' strategies in equilibrium, given their prior. Guerre et al. (2000) provide a technique to recover bidders' unobserved beliefs from the observed distribution of bids, which exploits the fact that the set of observed bids are the bidders' strategy profile in a pure-strategy equilibrium for a given realisation of signals. They show that, assuming equilibrium, the realized bid distribution can be used to estimate bidders' beliefs without specifying a functional form or making distributional assumptions.

Hortaçsu (2002) proposes a resampling procedure to implement this technique for multi-unit auctions, which consistently estimates each bidder's beliefs of the distribution of the auction prices, conditional on their strategy. Intuitively, each bidder is best responding to the strategies of other bidders, and each observed bid function corresponds to another bidder's optimal strategy for a particular signal realisation. And so, holding a bidder's strategy fixed, the observed bid distribution reflects the bidder's beliefs over the signal distribution. (In line with the model described in Section 3.1, a bidder who submits bids on both goods is treated as two separate bidders.)

The resampling procedure is as follows. Fixing bidder i in auction t, where i belongs to group $\mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, I construct a stratified sample of other bidders, denoted $\mathcal{O}_{i,t} = \{\mathcal{O}_{i,t}^{A,1}, \mathcal{O}_{i,t}^{A,2}, \mathcal{O}_{i,t}^{B,1}, \mathcal{O}_{i,t}^{B,2}\}$, from the bidding data as described below. The observation $(i', t') \in \mathcal{O}_{i,t}^{g,h}$ corresponds to the set of bids (i.e. individual bid function) submitted by bidder i' in group (g, h) in auction t' for a particular realisation of their signal. I draw $N^{g,h} - 1$ observations from $\mathcal{O}_{i,t}^{g,h}$ and $N^{g',h'}$ observations from $\mathcal{O}_{i,t}^{g',h'}, (g', h') \neq (g, h)$, with replacement. I calculate the equilibrium given the bids submitted by both the drawn bidders and the fixed bidder, and repeat this a large number of times for the fixed bidder. This yields an empirical distribution of auction prices, conditional on the bidder's own strategy.[26]

The resulting empirical distribution of auction prices for good g can be used to consistently estimate the expectation terms in Equation 2. The derivative term is estimated using a numerical derivative, as described in Kastl (2011). Given these estimates, the marginal values at the quantities demanded at the steps of bidder i's bid function can be recovered iteratively, starting with the value at step $k = K_i^g$ (for which $q_{i,k+1} = 0$).

The sample, $\mathcal{O}_{i,t}$, is a subset of all the bid functions submitted by the N potential bidders across the auctions. For the resampling procedure to yield a consistent estimator of bidder i's beliefs, two conditions on $\mathcal{O}_{i,t}$ must hold. First, the distribution of signals implied by $\mathcal{O}_{i,t}$ must be a consistent estimator for the true distribution of signals that bidder i faces in auction t. Second, the other bidders in $\mathcal{O}_{i,t}$ must face the same commonly known economic environment as bidder i faces in auction t and best respond to it. In the main analysis, I define $\mathcal{O}_{i,t} = \{(i', t') : (i', t') \neq (i, t), t' \in Y_t\}$, where Y_t is the set of auctions including auction t and its two nearest neighbors of the same term, and stratify by the grouping defined in Section 2.3.[27] This aims to maintain consistency in the economic environment and therefore signal distribution, while increasing precision by pooling bids from neighboring auctions.

4.3 Testing Procedure

Models $S(\rho)$ and T have different implications for bidders' marginal values and therefore for their bid shading. The testing procedure exploits these differences to compare the fit of the alternative models of behavior.

[26]Only a finite number of signal realisations is ever (implicitly) observed and the empirical distribution of auction prices may not have full support. This means that the conditional expectation terms may not be identified and the derivative of the auction price may be numerically unstable. So I use kernel density estimation to smooth the distribution and ensure full support, discretized at integer prices.

[27]The term to maturity of liquidity in each monthly auction follows a sequence. Two 3-month term auctions are followed by a 6-month term one. Anecdotal evidence suggests that some bidders participate in the ILTR auctions at regular 3- and 6-month intervals, with the intention to roll over the ILTR loan with the same piece of collateral. To eliminate the impact of correlated signals of this kind, I choose the two nearest neighbors, rather than three or more.

Ideally we would compare the bid shading implied by each candidate model to the actual amount of bid shading. Naturally, the candidate model with the smaller prediction error would be the better approximation.

Because the actual amount of bid shading is unobserved, I follow Backus et al.'s (2021) approach, which uses instrumental variables to mimic this comparison. Bidders' true marginal values for liquidity are assumed to be mean independent of the instruments. Under this assumption, bidders' values are uncorrelated with the projection of their bid shading on the instruments (i.e. the variation in the bid shading explained by the variation in the instruments). A particular candidate model is therefore a better approximation if it generates a smaller covariance between its estimated values and the projection of its estimated bid shading. Broadly speaking, the testing procedure is based on the relative size of these covariances.

The test compares the fit of two competing models of behavior. One advantage is that the test remains valid even if the two models are misspecified as it only assesses their relative performance.[28] This is particularly relevant in the empirical auction literature, in which a model is often required but is necessarily a simplification of reality.[29]

For a pair of models $m = \{1, 2\}$ and a measure of the model fit of each, $\mathcal{Q}_m$, the null hypothesis is that the two measures are equal, whereas the two alternatives are that the fit of one model is better than the other: $H_0 : \mathcal{Q}_1 = \mathcal{Q}_2, H_1 : \mathcal{Q}_1 < \mathcal{Q}_2,\ H_2 : \mathcal{Q}_1 > \mathcal{Q}_2$. In our test, a larger $\mathcal{Q}_m$ implies a worse model fit. So H_1 implies that Model 1 is a better fit than Model 2, and H_2 is analogous for Model 2.

The test is based on the relationship between the model-implied marginal values

[28] Grace (2024) analyses an external measure of fit of Model $S(0)$, by estimating lower bounds on the bidding costs implied by the model.

[29] For example, I compare Model $S(0)$, which assumes bidders have correct beliefs about the distribution of other bidders' bids, to Model T, in which bidders bid truthfully regardless of their beliefs. If neither model is correct, and the correct model is a version of Model $S(0)$, in which bidders have incorrect beliefs (or do not know the distribution of bids and their priors are incorrect), then the procedure tests whether Model T or Model $S(0)$ is a better approximation. For example, if bidders believe they face much stronger competition than in reality, Model T might be the better approximation.

and a function of a set of instruments, denoted $\boldsymbol{z}$, which are correlated with bidders' bid shading but uncorrelated with their true marginal values. One example of an instrument is the number of other bidders: with independent private values, the number of other bidders in the auction affects a bidder's market power but not their value. I define the true marginal value of bidder i at step k of their bid function in auction t by $v_{i,k,t}$, a function of a set of exogenous characteristics, $\boldsymbol{x}$, and an additively separable residual value, $\omega_{i,k,t}$. The critical assumption is that conditional on $\boldsymbol{x}$, the true residual value is mean independent of the instruments, $\boldsymbol{z}$, i.e. $\mathbb{E}\left[\omega_{i,k,t}|\boldsymbol{x},\boldsymbol{z}\right]=0$.

In each model, the bidder's marginal value is the sum of their observed bid and model-implied bid shading. Let $v^m_{i,k,t}=b_{i,k,t}+\mu^m_{i,k,t}$ be the marginal value of bidder i at step k in auction t under the assumptions of Model m, where $b_{i,k,t}$ is the bid price and $\mu^m_{i,k,t}$ is the bid shading. (The corresponding residual values implied by Model m are denoted $\omega^m_{i,k,t}$.) In Model T, $\mu^T_{i,k,t}=0$ so the bidder's value is simply their bid, $v^T_{i,k,t}=b_{i,k,t}$. By contrast, bid shading is not necessarily zero in Model $S(\rho)$, and the model-implied value is $v^{S(\rho)}_{i,k,t}=b_{i,k,t}+\mu^{S(\rho)}_{i,k,t}$.

To provide intuition for the test, I first show how an incorrect model is falsified in a simple example in which it is compared to the correct one. In this example, let $\boldsymbol{x}$ be a constant so that we can ignore the distinction between $v_{i,k,t}$ and $\omega_{i,k,t}$ and suppose that Model T is correct and Model $S(0)$ is incorrect. The bidder's true marginal value is therefore equal to the marginal value implied by Model T, $v_{i,k,t}=v^T_{i,k,t}=b_{i,k,t}$. Moreover, because their true value is equal to their bid, the marginal value implied by Model $S(0)$ is the sum of the true marginal value and the bid shading, i.e. $v^{S(0)}_{i,k,t}=b_{i,k,t}+\mu^{S(0)}_{i,k,t}=v_{i,k,t}+\mu^{S(0)}_{i,k,t}$. So the covariance between the model-implied marginal valuations and a valid instrument, e.g. $z_{i,k,t}$, will be zero for Model T and non-zero for Model $S(0)$: $\text{Cov}(v^T_{i,k,t},z_{i,k,t})=\text{Cov}(v_{i,k,t},z_{i,k,t})=0$ and $\text{Cov}(v^{S(0)}_{i,k,t},z_{i,k,t})=\text{Cov}(b_{i,k,t},z_{i,k,t})+\text{Cov}(\mu^{S(0)}_{i,k,t},z_{i,k,t})=\text{Cov}(v_{i,k,t},z_{i,k,t})+\text{Cov}(\mu^{S(0)}_{i,k,t},z_{i,k,t})=\text{Cov}(\mu^{S(0)}_{i,k,t},z_{i,k,t})\neq 0$. The crux of the test is that the correct model has zero covariance and a larger covariance implies a worse model fit, so the test is based on the relative magnitude of the covariances.

The steps of the test modified from Backus et al. (2021) are as follows.

1. For the two models $m = \{1, 2\}$,
   (a) For each step k of bidder i's bid function in auction t, estimate the bid shading, $\mu^m_{i,k,t}$, and calculate the implied values, $v^m_{i,k,t} = b_{i,k,t} + \mu^m_{i,k,t}$.
   (b) Estimate $v^m_{i,k,t}$ as a function of the set of exogenous characteristics $\boldsymbol{x}$ (specified below), $v^m_{i,k,t} = h(\boldsymbol{x}) + \omega^m_{i,k,t}$, and obtain the residuals, $\hat{\omega}^m_{i,k,t}$.
2. Estimate the difference in bid shading between the models as a function of $\boldsymbol{x}$ and the instruments, $\boldsymbol{z}$, i.e. $\Delta\mu_{i,k,t} = \mu^1_{i,k,t} - \mu^2_{i,k,t} = g(\boldsymbol{x}, \boldsymbol{z}) + \zeta_{i,k,t}$, where $\boldsymbol{z}$ and the function, $g(.)$, are specified below. Obtain the predictions, $\Delta\hat{\mu}_{i,k,t}$.[30]
3. For each model, compute the value of the moment, $\hat{Q}^m = \left(\frac{1}{n}\sum_{i,k,t}\hat{\omega}^m_{i,k,t}\Delta\hat{\mu}_{i,k,t}\right)^2$, where n is the number of observations.
4. Estimate the standard error, $\frac{\hat{\sigma}}{\sqrt{n}}$, of the difference between the moments, $(\hat{Q}^1 - \hat{Q}^2)$, by repeating Steps 1–3 on bootstrapped samples.[31]
5. Compute the test statistic, $T = \frac{\sqrt{n}(\hat{Q}^1 - \hat{Q}^2)}{\hat{\sigma}}$, which is distributed $\mathcal{N}(0, 1)$.

The set of exogenous characteristics, $\boldsymbol{x}$, which affect both bidders' values and their bid shading, includes auction-specific dummy variables, which control for heterogeneity in bidders' values across time, and a dummy variable equal to one if the bidder is large (as defined in Section 2.3.[32] The set $\boldsymbol{x}$ also includes a dummy variable equal to one if the bid corresponds to the last step in a bidder's bid function; this is intended to control for variation in bidders' values at the bid-bidder-auction level (our unit of

[30] This specification is chosen to maximize the power of the test (see Backus et al. (2021) for details).

[31] I follow Backus et al. (2021) and Roussille and Scuderi (2022) and treat the bid shading estimates for each model as data.

[32] Good-specific dummies are included in $\boldsymbol{x}$ in the test which pools bids across goods.

observation), conditional on the bidder's size and on the auction.[33,34]

For the set of instruments, $\boldsymbol{z}$, I use measures of the strength of competition that a bidder faces in the auction (listed below). The main assumption for these instruments to satisfy the exogeneity condition, i.e. $\mathbb{E}\left[\omega_{i,k,t}|\boldsymbol{x},\boldsymbol{z}\right] = 0$, is that bidders have independent private values (the IPV Assumption). Section 3.1 gives three reasons for why this assumption seems reasonable in my setting. For the instruments to be relevant, they must be correlated with the difference in bid shading between the competing models. Because one of the models is the conventional strategic model in which bidders are assumed to have consistent beliefs about the distribution of other bidders' bids, measures of the strength of competition are also relevant instruments.

There are many potential ways to measure the strength of competition. For a given bidder, this includes the number of bidders, or small or large bidders, for each good in the auction excluding the bidder themself (where size is defined in Section 2.3); the probability distribution of the auction price for each good that the bidder expects would occur if they did not bid at all, and moments of this distribution; and moments of the distribution of bids excluding the bidder's own bids. The IPV Assumption ensures that these statistics, which depend on other bidders' bids and their entry decisions, are valid instruments because bidder i's own valuation is mean independent of them.

Specifically, for bidder i in auction t, the set of instruments, which satisfy the exogeneity condition under the IPV Assumption, includes:

- For goods $g = \{A, B\}$, number of bidders, number of small bidders, and number of large bidders, each excluding bidder i, in auction t[35]
- The average bid in auction t, excluding the bids of bidder i

[33]The quantities demanded at the bids are excluded from $\boldsymbol{x}$ because, if Model $S(\rho)$ is the correct model of behavior, they are an equilibrium outcome and therefore endogenous to bidders' values. In principle, we could alternatively find additional instruments for the quantities demanded, or estimate bidders' marginal values at all quantities (not just the quantities demanded at the steps of the bidders' bid functions), which would require additional identifying conditions (see Grace (2024)).

[34]Online Appendix B repeats the analysis excluding the last-step dummy from $\boldsymbol{x}$.

[35]The model is estimated using a random forest so multicollinearity is not an issue.

- For $P^g_{i,t}$, calculated by excluding bidder i's bids when estimating the equilibrium distribution of the auction price (using the procedure in Section 4.2),
  - The conditional expected price for good A, $\mathbb{E}\left[P^A_{i,t} \leq X\right]$ for $X = [1, 5, 10]$
  - The conditional expected price for good B, $\mathbb{E}\left[P^B_{i,t} \leq X\right]$ for $X = [6, ..., 40]$
  - The conditional expected price differences for B: $(X - \mathbb{E}\left[X - 5 \leq P^B_{i,t} \leq X\right])$ for $X = [6, ..., 40]$ and $(X - \mathbb{E}\left[X - 10 \leq P^B_{i,t} \leq X\right])$ for $X = [6, ..., 40]$[36]
- All of the above set interacted with good-term-specific dummies

This implies a large number of potential instruments (equal to 460). Following Backus et al. (2021), I therefore project this set of instruments onto their principal components and select the three leading principal components as the three instruments to be used in the estimation; these explain 80% of the variation in the original to include in $\boldsymbol{z}$. The relationship between the instruments—measures of competition—and bid shading, are non-linear (see Equation 1 for Model $S(\rho)$). To capture these non-linearities, the functions $h(.)$ and $g(.)$ are estimated using a random forest.

I restrict the set of observations to bids strictly above the reserve prices.[37] Observations are weighted by the quantities demanded to ensure that the results are not driven by the fact that small bidders might submit bids far below their marginal values in Model $S(\rho)$ because the model implies that they are constrained in the number of bids that they submit.

5 Results

5.1 Model Fit

I run the test for pairs of models including Model T and Model $S(\rho)$ with different degrees of risk aversion. Results are reported in Table 4. The first set of results includes bidders for both goods. For these tests, Model $S(\rho)$ is estimated with a common

[36] The bounds of the conditional expectations represent sensible points on the distribution of bids (see Table 2). They capture the range of auction prices that are likely to occur.

[37] Bids at the reserve prices might correspond to bidders who would have optimally submitted bids below the reserve (corresponding to larger bid shading) in Model $S(\rho)$ if negative bids were allowed.

rationing coefficient for all bidders (written as a single parameter, ρ, for shorthand). The second and third sets of results only include bidders for good A and bidders for good B, respectively. For these cases, the test results for one good do not depend on the risk aversion of bidders for the other good, so I only specify the relevant risk aversion parameter (also written as a single parameter, ρ, for shorthand).

The test statistic when comparing Model T (Model 1) to Model $S(0)$ (Model 2) when pooling the two goods is -1.982, so Model T outperforms at the 2.5% significance level. Table 4 shows that this is driven by Model T providing significantly better fit than Model $S(0)$ for good B, but the models fitting equally well for good A. Moreover, Model T outperforms Model $S(\rho)$ for low levels of risk aversion and Model $S(\rho)$ fits the data better for higher values of ρ for good B, but the fit across models cannot be differentiated for good A. There are two possible explanations for these differences.

First, the definition of the two goods suggests that bidders for good B could plausibly be more risk averse than bidders for good A. Good B corresponds to funds lent against less liquid collateral assets, which typically have fairly illiquid secondary markets. A bidder for good B, who does not win in the BoE's auctions may have limited opportunities to obtain the loan elsewhere, unlike bidders for good A, who typically face a liquid secondary market.

Second, we have less power to discriminate between the different models for good A. 85.1% of bids for good A are between 0 and 5 basis points (and have to be a whole basis point), and the auction price is rarely above 2 basis points. This means that there is less variation in the observed bids, and therefore estimated values. This lack of variation makes it difficult to discriminate between models.

The results for good B show that Model T outperforms Model $S(\rho)$ for low degrees of risk aversion, suggesting that truthful bidding is a better approximation than the conventional risk-neutral strategic model. Moreover, Model $S(\rho)$ with high degrees of risk aversion outperforms Model $S(0)$, so risk aversion improves the fit of the

Table 4: Testing results

Model 1	Model 2						
Both goods							
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-1.982	-1.758	-1.662	-1.551	-1.487	-1.310	-1.085
$S(0)$		0.054	0.071	1.556	1.646	1.739	1.885
Good A							
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	0.215	0.024	0.024	0.024	0.024	0.024	0.193
$S(0)$		0.014	0.011	0.009	0.006	-0.002	-0.040
Good B							
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-2.574	-2.318	-2.112	-2.101	-1.882	-1.770	-1.447
$S(0)$		0.279	1.359	1.790	2.011	2.394	2.574

Note: Each entry shows the test statistic for Model 1 vs. Model 2, so that a negative entry indicates that Model 1 has better model fit than Model 2 (and a positive entry indicates the converse). The test statistic is distributed $\mathcal{N}(0,1)$ so the standard critical values apply: -1.645 for a 5% confidence level; -1.960 for 2.5%; -2.326 for 1%.

conventional model. Online Appendix B shows that the test results are very similar when Model $S(\rho)$ is estimated under the assumption that bidders do not tie.[38]

5.2 Rationalizing Truthful Bidding

The results of the pairwise tests show that if bidders are sufficiently risk averse, both Model T, i.e. truthful bidding, and Model $S(\rho)$, i.e. strategic bidding, fit the data equally well. This is because bidders' optimal strategies in both models are to bid approximately according to their true values if ρ is sufficiently high. Alternatively, if bidders have a low degree of risk aversion, Model T fits the data significantly better than Model $S(\rho)$. Moreover, truthful bidding is the most natural comparator to the conventional strategic model because bidders' optimal strategies in Model $S(\rho)$ may be to bid above or below their values. This suggests that a model in which bids correspond

[38] Appendix F of Grace (2024) also analyses an external measure of fit of Model $S(0)$ based on the size of the bidding costs required to explain the data under the assumptions of the model.

to bidders' true marginal values might be a reasonable approximation.

I consider two ways to rationalize this "truthful bidding". First, bidders are sufficiently risk averse that they bid truthfully even if they are best responding to correct beliefs about the economic environment and other participants' behavior, i.e. the assumptions of Model $S(\rho)$ hold. Second, regardless of their degree of risk aversion, bidders face sufficiently high costs of calculating the optimal strategy that it is more profitable to avoid this cost by bidding in a different way—simply bidding according to one's value is a natural alternative.[39] I consider these two explanations in turn.

Risk Aversion With a sufficiently high degree of risk aversion, Model $S(\rho)$ predicts that bidders will bid according to their true marginal values, i.e. Model $S(\rho)$ and Model T will predict the same behavior. This explains the results of the pairwise test in Table 4. The difference in model fit between Model T and Model $S(\rho)$ is insignificant if $\rho \geq 1$ (i.e. $\log(\rho) \geq 0$) for good B, and for all ρ for good A, implying that the models explain the data equally well.

There are two potential reasons why risk aversion may be a suitable assumption in the BoE's liquidity auctions. First, the incentives of the manager tasked with bidding on behalf of the participant may lead to risk-averse behavior within a principal-agent framework.[40] Second, the financial institutions themselves may be risk averse, given the auctions are intended to provide liquidity insurance, suggesting bidders may be willing to pay higher prices to reduce the uncertainty in their allocations. Because the sample period turned out to be relatively stable and participants' bids did not appear to reflect acute liquidity needs, the principal-agent interpretation seems more appropriate.

To examine the degree of risk aversion required for bidders to bid truthfully in Model $S(\rho)$, Figure 2 plots the absolute value of bid shading implied by Model $S(\rho)$ against the log of the risk aversion parameter, $\log(\rho)$. It shows that almost all bidders

[39]The two explanations are not mutually exclusive. The cost of calculating the optimal strategy required to rationalize truthful bidding is decreasing in the degree of risk aversion.

[40]Gordy (1999) describes a manager tasked with bidding in the auction, who has a concave payoff function either because of their own risk preferences or because of the remuneration structure.

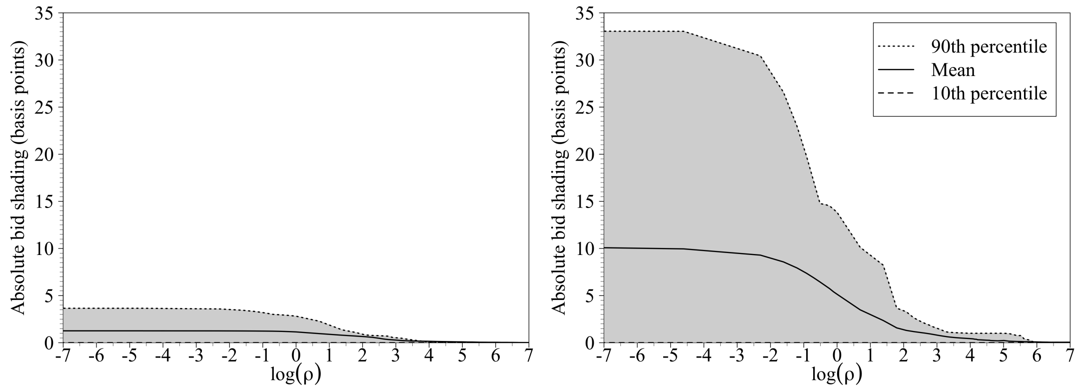

(a) Good A (b) Good B

Figure 2: Bid shading in Model $S(\rho)$

bid approximately truthfully in Model $S(\rho)$ if $\log(\rho) \geq 4$. In this range, the mean and 90th percentile of absolute bid shading for good B (weighted by quantity demanded) are less than 0.43 and 1 basis point, respectively. The mean and 90th percentile for good A are both less than 0.15 basis points.

To put this in context, a bidder in the 3-month auctions with $\log(\rho) = 4$ facing a 50-50 bet to lose £100,000 or gain X would accept if $X >$ £100,109.[41] (A bidder in the 6-month auctions would accept if $X >$ £100,209.)[42] This suggests that a fairly low degree of risk aversion is needed for truthful bidding to be optimal.

The degree of risk aversion that is sufficient for truthful bidding to be approximately optimal in Model $S(\rho)$ is consistent with the degree of risk aversion found in Treasury auctions, which are the closest settings to my study in which risk aversion has been estimated. Armantier and Sbaï (2006) estimate CARA parameters for small and large bidders in French Treasury auctions, which are approximately equivalent to $\log(\rho) = 9.83$ and $\log(\rho) = 5.04$, respectively. In US Treasury auctions, Boyarchenko et al. (2021) calibrate a parameter which is on average approximately equivalent to

[41] Online Appendix C provides the calculations for the results in this section.

[42] Aggregate bidder surplus, assuming bidders bid truthfully, is £220,461 per 3-month auction and £326,687 per 6-month auction on average (weighted by the quantities allocated per auction).

$\log(\rho) = 4.48$. Using any of these values in Model $S(\rho)$ in my setting would suggest that truthful bidding can be rationalized by the fact that bidders are risk averse.

However, the degree of risk aversion estimated in auctions varies across studies. When focusing only on primary dealers in Canadian Treasury auctions, Allen and Wittwer (2023) estimate a much smaller degree of risk aversion, approximately equivalent to $\log(\rho) = -4.57$. Unlike in Treasury auctions, bidders are not designated a "primary dealer" status in the BoE's liquidity auctions, and I do not have the necessary data to condition the risk aversion parameter on other bidder characteristics. However, it is plausible that risk aversion varies across bidders in my context. While my results suggest that risk aversion can explain truthful bidding when setting a common parameter ρ for all bidders, it may not be a sufficient explanation for individual bidders.

Cost of Sophistication An alternative way to explain truthful bidding is that a more sophisticated strategy is complicated and calculating it requires too many resources to be worthwhile. If the cost of calculating what would otherwise be the optimal strategy, i.e. the "cost of sophistication", is larger than the difference in expected utility between this strategy and bidding according to the bidder's true value for liquidity, the bidder prefers to bid truthfully. Hortaçsu and Puller (2008) find that a cost of sophistication best explains the fact that smaller bidders in Texan electricity markets persistently deviate from profit-maximizing behavior.

This interpretation of truthful bidding corresponds to Swinkels (2001) and Chakraborty and Englebrecht-Wiggans (2006), which show that the loss from bidding truthfully rather than strategically becomes arbitrarily small as the number of bidders increases in discrete multi-unit auctions. In these studies, truthful bidding is an ε-equilibrium, in which the "ε" can be interpreted as the cost of determining what otherwise would be the optimal strategy. If the incremental profit from this strategy is small, only a small cost of calculating it is required for truthful bidding to be a best response.

In the ILTR auctions, I estimate the incremental profit from what otherwise would be the optimal strategy for each bidder in the absence of a cost of sophistication, assuming bidders bid truthfully. If a bidder chooses to bids truthfully, the cost must be larger than the incremental profit. The estimate therefore provides a lower bound on the cost of sophistication that can explain why bidders bid truthfully.

Specifically, I compare the expected surplus for each bidder from submitting the optimal bid function with two steps, to submitting a bid function which corresponds to their true marginal valuation function, under the assumption that bidders are truthful and risk neutral.[43,44] The expected surplus from the optimal strategy would be larger if we estimated the optimal bid function with three or more steps rather than two (which would increase the lower bound), but it is likely that the difference would be small.[45]

I find that the average lower bound on the cost of sophistication is 4.28% of bidder surplus for bidders for good A and 5.11% for bidders for good B. This suggests that a relatively small cost of determining the optimal strategy can rationalize truthful bidding. Equivalently, if bidders' actual strategies are to bid truthfully, they obtain up to 95% of the surplus that would have been generated by the two-step optimal bid function. This suggests that a relatively small amount of surplus is 'left on the table' by bidding truthfully. By comparison, Hortaçsu and Puller (2008) find that bidders' actual strategies generate between 0% and 80% (excluding loss-making bidders) of that generated by the optimal response (based on the information available to the bidders at the time of bidding).

[43] I estimate the surpluses and optimal two-step bid function by estimating the distribution of other bidders' bids and implementing a grid search across all bid prices and quantity increments of 0.25% of the maximum supply. For bidders whose bid functions do not lie on this grid, I take the lower envelope of their bid functions as their "truthful" bid function (for bidders who submit more than two bids, I take the lower envelope of the highest two bids they submit). This approach is described in Grace (2024).

[44] Risk aversion would reduce the incremental profit from the optimal strategy.

[45] One reason is that Grace (2024) estimates the surpluses of optimal bid functions with different numbers of steps and finds that the benefit of a second step can be large whereas the benefit of a third step is small. Moreover, Kastl (2012) shows that the loss from using a step function rather than a continuous bid function decreases at a quadratic rate, suggesting that the additional surplus from additional steps would also be small. Estimating the optimal bid function becomes very computationally intensive with many steps, so I only estimate the two-step optimal bid function.

6 Conclusion

The Bank of England (BoE) introduced its liquidity auctions both to efficiently provide liquidity insurance to financial institutions and to provide better information to the BoE about market conditions. The design—a Product-Mix Auction—is efficient if bids correspond to bidders' true marginal values for loans. Moreover, Paul Fisher (then Executive Director at the BoE) noted at the time that the pattern of bids gives a signal of market stress "because bids in the auctions should provide accurate information on individual banks' demand for liquidity and the prices they are willing to pay for it" (Fisher, 2011a pp.11). That is, bids were understood as expressing bidders' true valuations for loans. If the bids were instead viewed through the model of behavior which is more conventional in the literature—bidders choose to submit bids which may differ from their true values in order to maximize their own expected surplus—the measured efficiency of, and information gleaned from, the auctions would change.

I therefore compare the relative performance of these alternative models of behavior. I find that bidding behavior is better explained by a model of "truthful bidding" than by a conventional model in which bidders are both strategic and risk neutral, using a testing procedure based on model fit developed by Backus et al. (2021). Moreover, I build on Kastl's (2011) framework to allow for risk aversion, and I find that the degree of risk aversion required for truthful bidding to be approximately optimal within the conventional model is consistent with that found in studies of risk aversion that are the most similar to my setting. Alternatively, truthful bidding can be rationalized by a cost of determining the optimal strategy, and so can be interpreted as an ε-equilibrium as in Swinkels (2001) and Chakraborty and Englebrecht-Wiggans (2006).

Importantly, I have not confirmed that bidders do indeed bid truthfully in the BoE's liquidity auctions but only that this simple model of behavior—truthful bidding—is an improvement on the conventional model. There may be alternative models that I have not considered that fit the data even better.

References

Allen, J., and M. Wittwer. 2023. "Intermediary Asset Pricing: Capital Constraints and Market Power." Bank of Canada Staff Working Paper 2023-51.

Armantier, O., and E. Sbaï. 2006. "Estimation and Comparison of Treasury Auction Formats when Bidders are Asymmetric." *Journal of Applied Econometrics*, 21(6): 745-779.

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

Backus, M., C. Conlon, and M. Sinkinson. 2021. "Common ownership and competition in the ready-to-eat cereal industry." NBER Working Paper No. w28350.

Bank of England. 2010a. "Bank of England Market Notice: Indexed Long-Term Repo Operations; US Dollar Repo Operations." 26 May.

Bank of England. 2010b. *Quarterly Bulletin*, 2, Bank of England.

Bank of England. 2011. *Quarterly Bulletin*, 1, Bank of England.

Berry, S., and P. Haile. 2014. "Identification in differentiated products markets using market level data." *Econometrica*, 82(5): 1749-1797.

Borenstein, S., J. B. Bushnell, and F. A. Wolak. 2002. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review*, 92: 1376–1405.

Boyarchenko, N., D. O. Lucca, and L. Veldkamp. 2021. "Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Auctions." *Journal of Political Economy*, 129(2): 607-645.

Cassola, N., A. Hortaçsu, and J. Kastl. 2013. "The 2007 Subprime Market Crisis through the Lens of European Central Bank Repo Auctions for Short-Term Funds." *Econometrica*, 81(4): 1309–45.

Chakraborty, I., and R. Engelbrecht-Wiggans. 2014. "Truthful Bidding in Uniform-Price Multi-Unit Auctions." Unpublished.

Chapman, J. T. E., D. McAdams, and H. J. Paarsch. 2006. “Bounding Best-Response Violations in Discriminatory Auctions with Private Values.” Unpublished.

de Roure, C., and N. McLaren. 2021. “Liquidity Transformation, Collateral Assets and Counterparties.” *Central Bank Review*, 21(4): 119-129.

Duarte, M., L. Magnolfi, and C. Roncoroni. 2023. “The competitive conduct of consumer cooperatives.” *RAND Journal of Economics*, forthcoming.

Fisher, P. 2011a, March. “Recent developments in the sterling monetary framework.” Lecture given at Manchester Economics Seminar, University of Manchester.

Fisher, P. 2011b, April. “Central Bank Policy on Collateral.” Bank of England.

Friedman, M. 1960. *A Program for Monetary Stability*. New York: Fordham University Press.

Friedman, M. 1991, August 28. “How to sell government securities.” *Wall Street Journal*.

Giese, J., and C. Grace. 2023. “An evaluation of the Bank of England’s ILTR operations: comparing the product-mix auction to alternatives.” Bank of England Staff Working Paper, No. 1,044.

Gordy, M. B. 1999. “Hedging Winner’s Curse with Multiple Bids: Evidence from the Portuguese Treasury Bill Auction.” *Review of Economics and Statistics*, 81(3): 448-465.

Grace, C. 2024. “Competing models of the Bank of England’s liquidity auctions: truthful bidding is a good approximation.” Bank of England Staff Working Paper, No. 1,061.

Guerre, E., I. Perrigne, and Q. Vuong. 2000. “Optimal Nonparametric Estimation of First-Price Auctions.” *Econometrica*, 68(3): 525–74.

Gupta, S., and R. Lamba. 2023. “Treasury auctions during a crisis.” https://www.rohitlamba.com/. Accessed 2023-11-01.

Häfner, S. 2023. “Risk Aversion in Share Auctions: Estimating Import Rents from

TRQs in Switzerland." *Quantitative Economics*, 14(2): 419-470.

Haile, P., H. Hong, and M. Shum. 2003. "Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions." NBER Working Paper No. w10105.

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

Hortaçsu, A., and J. Kastl. 2012. "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions." *Econometrica*, 80(6): 2511–42.

Hortaçsu, A., J. Kastl, and A. Zhang. 2018. "Bid Shading and Bidder Surplus in the US Treasury Auction System." *American Economic Review*, 108(1): 147–69.

Hortaçsu, A., F. Luco, S. L. Puller, and D. Zhu. 2019. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." *American Economic Review*, 109(12): 4302–4342.

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): 833–65.

Hortaçsu, A., and D. McAdams. 2018. "Empirical Work on Auctions of Multiple Objects." *Journal of Economic Literature*, 56(1): 157-84.

Hortaçsu, A., and I. Perrigne. 2021. "Empirical perspectives on auctions." In *Handbook of Industrial Organization*, vol. 5, 1: 81-175. Elsevier.

Hortaçsu, A., and S. L. Puller. 2008. "Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market." *RAND Journal of Economics*, 39(1): 86–114.

Kastl, J. 2011. "Discrete Bids and Empirical Inference in Divisible Good Auctions." *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): 339–52.

Klemperer, P. 2008. “A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions.” Oxford University.

Klemperer, P. 2010. “The Product-Mix Auction: a New Auction Design for Differentiated Goods.” *Journal of the European Economic Association*, 8(2-3): 526-36.

Klemperer, P. 2018. “Product-Mix Auctions.” Working Paper 2018-W07, Nuffield College.

Klemperer, P., and M. Meyer. 1989. “Supply Function Equilibria in Oligopoly under Uncertainty.” *Econometrica*, 57(6): 1243–77.

Krishna, V. 2009. *Auction theory.* Hoboken, NJ: Academic Press.

McAdams, D. 2008. “Partial Identification and Testable Restrictions in Multi-Unit Auctions.” *Journal of Econometrics*, 146(1): 74–85.

Rivers, D., and Q. Vuong. 2002. “Model selection tests for nonlinear dynamic models.” *The Econometrics Journal*, 5(1): 1-39.

Roussille, N., and B. Scuderi. 2021.“ Bidding for Talent: Equilibrium Wage Dispersion on a High-Wage Online Job Board.” Unpublished.

Starc, A., and T. G. Wollmann. 2022. “Does Entry Remedy Collusion? Evidence from the Generic Prescription Drug Cartel.” NBER Working Paper No. w29886.

Swinkels, J. M. 2001. “Efficiency of large private value auctions.” *Econometrica*, 69(1): 37-68.

Wilson, R. 1979. “Auctions of Shares.” *Quarterly Journal of Economics*, 93(4): 675–89.

Winters, B. 2012. “Review of the Bank of England’s Framework for Providing Liquidity to the Banking System.” Bank of England.

Wolak, F. A. 2007. “Quantifying the supply-side benefits from forward contracting in wholesale electricity markets.” *Journal of Applied Econometrics*, 22(7): 1179-1209.

Wolfram, C. D. 1999. “Measuring Duopoly Power in the British Electricity Spot Market.” *American Economic Review*, 89(4): 805–826.

For Online Publication

A Mathematical Appendix

A.1 Proof of Proposition 1

Kastl (2011) proves Proposition 1 for $\rho^g = 0$. I follow his approach for the case of $\rho^g > 0$. With a slight abuse of notation, it is helpful to define $P^g(0)$ and $y_i(0)$ by the equilibrium auction price of good g and bidder i's equilibrium allocation, respectively, under the original strategy profile, $\boldsymbol{\sigma}(\boldsymbol{\Theta})$; and to define $P^g(\epsilon)$ and $y_i(\epsilon)$ by the equilibrium auction price of good g and bidder i's equilibrium allocation, respectively, if bidder i unilaterally deviates to a strategy in which their type θ_i demands a quantity $(q_{i,k} - \epsilon)$, $\epsilon > 0$, at step k and the rest of the strategy profile, including the rest of their strategy, is unchanged.

For $\rho^g > 0$, we aim to evaluate the limit $\lim_{\epsilon \to 0} \frac{1}{\epsilon} \left(\mathbb{E}\left[\frac{1}{\rho^g} \left(1 - e^{-\rho^g U_i(y_i(\epsilon), P^g(\epsilon))}\right)\right] - \mathbb{E}\left[\frac{1}{\rho^g} \left(1 - e^{-\rho^g U_i(y_i(0), P^g(0))}\right)\right]\right)$ where $U_i(.)$ is shorthand for $U_i(.|\theta_i)$ and $\mathbb{E}[.]$ is shorthand for $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}[.]$. Equation 1 divides the impact of bidder i's deviation into the impact on the quantity allocated to bidder i and the impact on the expected auction price conditional on winning. To do this, I follow Kastl's (2011) approach by partitioning the state space as follows:

$$\theta_{1k}(x) = \{\boldsymbol{\Theta_{-i}} : b_{i,k+1} < P^g(x) \leq b_{i,k},\ y_i(x) = q_{i,k} - x\}$$

$$\theta_{2k}(x) = \left\{\boldsymbol{\Theta_{-i}} : P^g(x) = b_{i,k},\ y_i(x) = y_i^R(q_{i,k} - x - q_{i,k-1}) : q_{i,k-1} < y_i^R < q_{i,k} - x\right\}$$

$$\theta_{3k}(x) = \left\{\boldsymbol{\Theta_{-i}} : P^g(x) = b_{i,k+1},\ y_i(x) = y_i^R(q_{i,k+1} - q_{i,k} + x) : q_{i,k} - x < y_i^R < q_{i,k+1}\right\}$$

$$\theta_{4k}(x) = \{\boldsymbol{\Theta_{-i}} : b_{i,k} < P^g(x),\ y_i(x) \leq q_{i,k-1}\}$$

$$\theta_{5k}(x) = \{\boldsymbol{\Theta_{-i}} : P^g(x) < b_{i,k+1},\ q_{i,k+1} \leq y_i(x)\}$$

for $x \in \{0, \epsilon\}$, where $\epsilon > 0$ and $y_i^R(q)$ is the quantity that bidder i is allocated when rationed, given the quantity they demand at the margin is q. We then define subsets of θ_{2k} and θ_{3k}: $\omega_{1k}(\epsilon)\,\{\boldsymbol{\Theta_{-i}} : \boldsymbol{\Theta_{-i}} \in \theta_{2k}(0) \cap \theta_{1k}(\epsilon)\}$, $\omega_{2k}(\epsilon) = \{\boldsymbol{\Theta_{-i}} : \boldsymbol{\Theta_{-i}} \in \theta_{2k}(0) \cap \theta_{3k}(\epsilon)\}$, and $\omega_{3k}(\epsilon) = \{\boldsymbol{\Theta_{-i}} : \boldsymbol{\Theta_{-i}} \in \theta_{1k}(0) \cap \theta_{3k}(\epsilon)\}$.

We have that $\sum_{j=1}^{3}\mathbb{P}\left(\theta_{jk}(0)\right)=\sum_{j=1}^{3}\mathbb{P}\left(\theta_{jk}(\epsilon)\right)$ because the deviation only changes bidder i's outcome in cases in which the auction price is weakly between their bid prices at steps k and $k+1$. By the Law of Total Probability, we can evaluate the utility function piecewise using the definition of the partition. We therefore aim to evaluate:

$$
\begin{aligned}
&\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i(\epsilon),P^g(\epsilon))}\Big)\Big]-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i(0),P^g(0))}\Big)\Big]\Big)\\
&=\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\Big);\theta_{1k}(0)\Big]-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\Big);\theta_{1k}(0)\Big]\Big)\\
&+\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\big);\theta_{2k}(0)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-q_{i,k-1}),b_{i,k})}\big);\theta_{2k}(0)\Big]\Big)\\
&+\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\big);\theta_{3k}(0)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\big);\theta_{3k}(0)\Big]\Big)\\
&+\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\big)-\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},P^g(0))}\big);\cup_{j=1}^{3}\theta_{jk}(0)\Big]\Big)\\
&+\lim_{\epsilon\to0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k},b_{i,k})}\Big);\omega_{1k}(\epsilon)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\big);\omega_{1k}(\epsilon)\Big]\\
&+\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},b_{i,k})}\big);\omega_{2k}(\epsilon)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\big);\omega_{2k}(\epsilon)\Big]\\
&+\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\big);\omega_{3k}(\epsilon)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},b_{i,k+1})}\big);\omega_{3k}(\epsilon)\Big]\\
&+\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\big);\omega_{2k}(\epsilon)\Big]\\
&\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},b_{i,k+1})}\big);\omega_{2k}(\epsilon)\Big]\\
&+\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\big);\omega_{1k}(\epsilon)\Big]-\mathbb{E}\Big[\tfrac{1}{\rho^g}\big(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\big);\omega_{1k}(\epsilon)\Big]\\
&-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\Big);\omega_{3k}(\epsilon)\Big]+\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\Big);\omega_{3k}(\epsilon)\Big]\Big)
\end{aligned}
$$

It follows from the definition of $U_i(y_i(x),P^g(x))$ that $\frac{\partial U_i(y_i(x),P^g(x))}{y_i(x)}=v^{g,h}(y_i(x),\theta_i)-$

$P^g(x)$. Applying L'Hôpital's rule, the first three limits are

$$\lim_{\epsilon\to 0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\Big);\theta_{1k}(0)\Big]-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\Big);\theta_{1k}(0)\Big]\Big)$$
$$=\mathbb{E}\Big[e^{-\rho^g U_i(q_{i,k},P^g(0))}\Big(v^{g,h}(q_{i,k},\theta_i)-P^g(0)\Big);\theta_{1k}(0)\Big],$$

$$\begin{aligned}&\lim_{\epsilon\to 0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\Big);\theta_{2k}(0)\Big]\\&\qquad\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i^R(q_{i,k}-q_{i,k-1}),b_{i,k})}\Big);\theta_{2k}(0)\Big]\Big)\\&=\mathbb{E}\Big[e^{-\rho^g U_i(y_i^R(q_{i,k}-q_{i,k-1}),b_{i,k})}\Big(v^{g,h}(y_i^R(q_{i,k}-q_{i,k-1}),\theta_i)-b_{i,k}\Big);\theta_{2k}(0)\Big],\end{aligned}$$

and

$$\begin{aligned}&\lim_{\epsilon\to 0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\Big);\theta_{3k}(0)\Big]\\&\qquad\qquad-\mathbb{E}\Big[\tfrac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\Big);\theta_{3k}(0)\Big]\Big)\\&=\mathbb{E}\Big[e^{-\rho^g U_i(y_i^R(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\Big(v^{g,h}(y_i^R(q_{i,k+1}-q_{i,k}),\theta_i)-b_{i,k+1}\Big);\theta_{3k}(0)\Big]\end{aligned}$$

Now consider the fourth limit. We can show that the expectation term is continuous in ϵ by first partitioning $\mathbb{E}\Big[\frac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\right);\cup_{j=1}^3\theta_{jk}(0)\Big]$, into $\mathbb{E}\Big[\frac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\right)\Big|\theta_{1k}(0)\Big]\mathbb{P}\left(\theta_{1k}(0)\right)$, $\mathbb{E}\Big[\frac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},b_{i,k})}\right)\Big|\theta_{2k}(0)\Big]\mathbb{P}\left(\theta_{2k}(0)\right)$, and $\mathbb{E}\Big[\frac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},b_{i,k+1}}\right)\Big|\theta_{3k}(0)\Big]\mathbb{P}\left(\theta_{3k}(0)\right)$. Because both $\left[P^g(\epsilon)\middle|\theta_{1k}(0)\right]\in[b_{i,k+1},b_{i,k}]$ (so that $\left[U_i(q_{i,k},P^g(\epsilon))\middle|\theta_{1k}(0)\right]\geq 0$) and $\rho^g>0$, it follows that the exponential is bounded, i.e. $\left[e^{\rho^g U_i(q_{i,k},P^g(\epsilon))}\middle|\theta_{1k}(0)\right]\in[0,1]$. Therefore by the same reasoning as Kastl's (2011) Lemmas A3 and A4, $\mathbb{E}\left[\frac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\right)\middle|\theta_{1k}(0)\right]$ is continuous in ϵ at $\epsilon=0$ for a.e. $\theta_i\in\Theta_i$, is of bounded variation and satisfies the Luzin N property and so is locally differentiable with respect to ϵ at $\epsilon=0$ for a.e. $\theta_i\in\Theta_i$.

And so, applying L'Hôpital's rule to the fourth limit:

$$\begin{aligned}&\lim_{\epsilon\to 0}\tfrac{1}{\epsilon}\Big(\mathbb{E}\left[\tfrac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\right)-\tfrac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(0))}\right);\cup_{j=1}^{3}\theta_{jk}(0)\right]\Big)\\&\quad=\tfrac{\partial}{\partial\epsilon}\Big(\mathbb{E}\left[\tfrac{1}{\rho^g}\left(1-e^{-\rho^g U_i(q_{i,k},P^g(\epsilon))}\right);\cup_{j=1}^{3}\theta_{jk}(0)\right]\Big)\Big|_{\epsilon=0}\end{aligned}$$

And finally, we can express the expectations in the last limit in terms of their conditional expectations, i.e. $\mathbb{E}\Big[\frac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y^g_{i,k}(x),P^g(x))}\Big);\omega_{jk}(\epsilon)\Big]=\mathbb{E}\Big[\frac{1}{\rho^g}\Big(1-e^{-\rho^g U_i(y^g_{i,k}(x),P^g(x))}\Big)\Big|\omega_{jk}(\epsilon)\Big]\mathbb{P}\big(\omega_{jk}(\epsilon)\big)$ for $j\in\{1,2,3\}$ and note that $\lim_{\epsilon\to 0}\big(y_i^R(q_{i,k}-\epsilon-q_{i,k-1})\big|\omega_{jk}(\epsilon)\big)=q_{i,k}$; $\lim_{\epsilon\to 0}\big(y_i^R(q_{i,k+1}-q_{i,k-1}+\epsilon)\big|\omega_{jk}(\epsilon)\big)=q_{i,k}$; $\lim_{\epsilon\to 0}\big(q_{i,k-1}-\epsilon\big|\omega_{jk}(\epsilon)\big)=q_{i,k}$; and, because Kastl's (2011) Lemma 1 holds, $\lim_{\epsilon\to 0}\big(\mathbb{P}\big(\omega_{jk}(\epsilon)\big)\big)=0$. It follows from L'Hôpital's rule that the last limit is zero (see Kastl (2011) for details).

Combining these results, recalling the definitions of the θ_{jk}s, observing both that $\mathbb{I}(b_{i,k}\geq P^g(\epsilon)\geq b_{i,k+1})=\mathbb{I}\left(b_{i,k}\geq P^g(0)\geq b_{i,k+1}\right)$ and that the deviation only affects the quantity allocated to bidder i in the event of being rationed (i.e. in the sets θ_{2k} and θ_{3k}) if another bidder is also rationed (i.e. bidder i "ties" on the margin), the necessary condition for equilibrium, for $\rho^g>0$, is the second case of Equation 1.

A.2 Equation 1 in the Case of No Ties

In the case that there are no ties on the margin for bidder i, the condition which rules out profitable local deviations in the quantity that they demand is particularly simple. The bidder's gross utility, $\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du$, is deterministic, conditional on winning, as are their bidding costs, $c_i(K_i)$, so these elements do not enter into the bidder's tradeoff, even if they are risk averse. The only random component in their utility function is the price that they must pay for the units that they win. If there are no ties then $\mathbb{P}(P^g = b_{i,k+1} \wedge Tie^g) = \mathbb{P}(P^g = b_{i,k} \wedge Tie^g) = 0$, Equation 1 becomes

$$\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho^g U_i(q_{i,k},P^g|\theta_i)}\left(v^{g,h}(q_{i,k},\theta_i) - P^g\right)\middle|b_{i,k} > P^g > b_{i,k+1}\right]\mathbb{P}(b_{i,k} > P^g > b_{i,k+1})$$

$$= \begin{cases} q_{i,k}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g = 0 \\ \frac{1}{\rho^g}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho^g U_i(q_{i,k},\tilde{P}^g(\epsilon)|\theta_i)}\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g > 0 \end{cases}$$

$$\Longrightarrow \mu_{i,k} = v^{g,h}(q_{i,k},\theta_i) - b_{i,k}$$

$$= \frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho^g P^g q_{i,k}}P^g\middle|b_{i,k} > P^g > b_{i,k+1}\right]}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho^g P^g q_{i,k}}\middle|b_{i,k} > P^g > b_{i,k+1}\right]} - b_{i,k}$$

$$+ \begin{cases} \Omega q_{i,k}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g = 0 \\ \Omega\frac{\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho^g \tilde{P}^g(\epsilon)q_{i,k}}\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0}}{\rho^g\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho^g P^g q_{i,k}}\middle|b_{i,k} > P^g > b_{i,k+1}\right]} & \text{if } \rho^g > 0 \end{cases} \tag{3}$$

where $\Omega = \frac{1}{\mathbb{P}(b_{i,k} > P^g > b_{i,k+1})}$.

A.3 Derivation of Equation 2

The main complexity in identfying bidders' marginal values in Model $S(\rho)$ comes from allowing for ties. If both a bidder's bid equals the auction price and another bidder submits a bid at that same price, a marginal increase in the quantity the bidder demands at step k increases the quantity they are allocated by the product of the marginal quantity they demand at step k and the rationing coefficient.[46] The bidder's deviation in the quantity they demand therefore affects their utility differently in the case that their bid is marginal relative to the case that their bid is strictly above the auction price. As described in Section 4.1, I therefore make two additional assumptions to identify bidders' values:

[46] The bidder is allocated the entire quantity they demand at step $k-1$ as well as an amount which is proportional to the marginal quantity they demand at step k. If the bidder ties, the deviation increases their allocation by the product of the marginal quantity they demand at step k and the rationing coefficient because bids are rationed pro-rata. If the bidder does not tie, they are allocated the entire quantity supplied on the margin, so a deviation in the quantity they demand has no impact on their allocation.

Assumption 1 *For bidder $i \in \mathcal{N}^{g,h}$, $g \in \{A, B\}$, $h \in \{1, 2\}$, $v^{g,h}(y_i, \theta_i) = v^{g,h}(q_{i,k}, \theta_i)$ $\forall y_i \in \Big[q_{i,k-1} + r^g_{i,k}(q_{i,k} - q_{i,k-1}), q_{i,k} + r^g_{i,k+1}(q_{i,k+1} - q_{i,k})\Big)$.*

Assumption 2 *The rationing coefficient for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}$ at step $k \in \{1, ..., K_i\}$, conditional on the strategy profile $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, denoted $r^g_{i,k}$, is deterministic.*

Under Assumptions 1 and 2, Equation 1 simplifies as follows. (Alternative natural assumptions would achieve the same result.) First note that Assumption 2 implies that $\mathbb{E}\left[\frac{\partial y_i}{\partial q_{i,k}}\middle| P^g = b_{i,k} \wedge Tie^g\right] = r^g_{i,k}$ and $\mathbb{E}\left[\frac{\partial y_i}{\partial q_{i,k}}\middle| P^g = b_{i,k+1} \wedge Tie^g\right] = (1 - r^g_{i,k+1})$. Splitting up $U_i(q_{i,k}, P^g|\theta_i)$ and $U_i(q_{i,k}, \tilde{P}^g(\epsilon)|\theta_i)$ into their constituent parts and cancelling the deterministic components of the bidder's expected utility implies

$$
\begin{aligned}
&\mathbb{E}\Big[e^{\rho^g P^g q_{i,k}}\left(v^{g,h}(q_{i,k}, \theta_i) - P^g\right)\Big| P^g \in (b_{i,k+1}, b_{i,k})\Big]\mathbb{P}(P^g \in (b_{i,k+1}, b_{i,k}))\\
&+ \Big(e^{\rho^g \int_{(q_{i,k-1}+r^g_{i,k}(q_{i,k}-q_{i,k-1}))}^{q_{i,k}} v^{g,h}(u,\theta_i)du} e^{\rho^g b_{i,k}\left(q_{i,k-1}+r^g_{i,k}\left(q_{i,k}-q_{i,k-1}\right)\right)}\\
&\quad \left(v^{g,h}\left(q_{i,k-1} + r^g_{i,k}\left(q_{i,k} - q_{i,k-1}\right), \theta_i\right) - b_{i,k}\right) r^g_{i,k}\Big)\mathbb{P}(P^g = b_{i,k} \wedge Tie^g)\\
&+ \Big(e^{-\rho^g \int_{q_{i,k}}^{(q_{i,k}+r^g_{i,k+1}(q_{i,k+1}-q_{i,k}))} v^{g,h}(u,\theta_i)du} e^{\rho^g b_{i,k+1}\left(q_{i,k}+r^g_{i,k+1}\left(q_{i,k+1}-q_{i,k}\right)\right)}\\
&\quad \left(v^{g,h}\left(q_{i,k} + r^g_{i,k+1}\left(q_{i,k+1} - q_{i,k}\right), \theta_i\right) - b_{i,k+1}\right)\left(1 - r^g_{i,k+1}\right)\Big)\mathbb{P}(P^g = b_{i,k+1} \wedge Tie^g)\\
&= \begin{cases} q_{i,k}\frac{\partial}{\partial \epsilon}\left(\mathbb{E}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(\tilde{P}^g(\epsilon) \in [b_{i,k+1}, b_{i,k}]\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g = 0\\ \frac{1}{\rho^g}\frac{\partial}{\partial \epsilon}\left(\mathbb{E}\left[e^{\rho^g \tilde{P}^g(\epsilon) q_{i,k}}\mathbb{I}\left(\tilde{P}^g(\epsilon) \in [b_{i,k+1}, b_{i,k}]\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho^g > 0 \end{cases}
\end{aligned}
$$

Then, applying Assumption 1 to the marginal valuation function yields Equation 2.

Table 5: Testing results assuming no ties

Both goods							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-1.764	-1.701	-1.655	-1.456	-1.408	-1.375	-1.243
$S(0)$		0.064	0.879	1.145	1.281	1.428	1.670
Good A							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	0.019	0.018	0.018	0.017	0.017	0.017	0.014
$S(0)$		0.004	0.004	0.004	0.004	0.004	-0.046
Good B							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-2.322	-2.223	-2.174	-2.094	-2.014	-1.941	-1.770
$S(0)$		0.053	0.088	0.120	0.149	1.673	2.125

Note: Each entry shows the test statistic for Model 1 vs. Model 2, so that a negative entry indicates that Model 1 has better model fit than Model 2 (and a positive entry indicates the converse). The test statistic is distributed $\mathcal{N}(0, 1)$ so the standard critical values apply: -1.645 for a 5% confidence level; -1.960 for 2.5%; -2.326 for 1%.

B Robustness of the Test Results

Table 5 shows the test results with the bid shading in Model $S(\rho)$ estimated under the assumption that bidders do not tie and the testing specification identical to the main analysis. Figure 3 plots the absolute value of bid shading implied by Model $S(\rho)$ against the log of the risk aversion parameter, $\log(\rho)$, assuming bidders do not tie.

Table 6 shows the test results, with the bid shading in Model $S(\rho)$ estimated under the assumption of the main model and the testing specification identical to the main analysis except that the dummy variable for a last step in a bidder's bid function is excluded from the set of characteristics, $\boldsymbol{x}$.

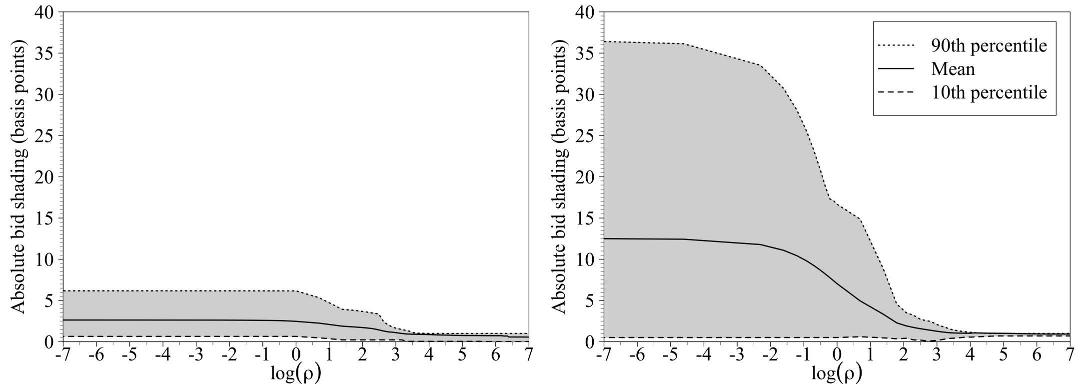

(a) Good A (b) Good B

Figure 3: Bid shading in Model $S(\rho)$ assuming no ties

Table 6: Testing results excluding the dummy for last step from $\boldsymbol{x}$

Both goods							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-1.609	-1.511	-1.464	-1.307	1.295	-1.101	-1.000
$S(0)$		0.060	0.085	0.094	0.099	1.342	1.724
Good A							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	0.024	0.024	0.025	0.025	0.024	0.023	0.023
$S(0)$		0.005	0.006	0.006	0.005	0.004	-0.024
Good B							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-2.144	-2.032	-1.952	-1.875	-1.798	-1.739	-1.494
$S(0)$		0.137	0.152	0.175	1.770	1.875	2.103

Note: Each entry shows the test statistic for Model 1 vs. Model 2, so that a negative entry indicates that Model 1 has better model fit than Model 2 (and a positive entry indicates the converse). The test statistic is distributed $\mathcal{N}(0,1)$ so the standard critical values apply: -1.645 for a 5% confidence level; -1.960 for 2.5%; -2.326 for 1%.

C Calculations for Section 5.2

The risk aversion parameter, ρ, measures the rate at which a bidder's marginal utility decreases when their net utility from the auction increases by one unit of the maximum supply. Maximum supply is £5 billion in the 3-month term auctions and £2.5 billion in the 6-month term auctions so the risk aversion parameters measured in £ are $\rho^M = (0.2 \times 10^{-9} \times \rho)$ and $\rho^M = (0.4 \times 10^{-9} \times \rho)$, respectively.

Let W be the bidder's initial net utility. A bidder will be indifferent to the 50-50 bet to lose Y and win X if $\frac{1}{2}\frac{1}{\rho^M}\left(1 - \exp^{-\rho^M(W+X)}\right) + \frac{1}{2}\frac{1}{\rho^M}\left(1 - \exp^{-\rho^M(W-Y)}\right) = \frac{1}{\rho^M}\left(1 - \exp^{-\rho^M(W)}\right)$. And so, if $Y = 100,000$ and $\log(\rho) = 4$ for a bidder in the 3-month auctions, so that $\rho^M = 1.092 \times 10^{-8}$, then $X =$ £100,109. For a bidder in the 6-month auctions, $\rho^M = 2.18 \times 10^{-8}$ so $X =$ £100,209.

Armantier and Sbaï (2006) estimate CARA parameters of 6.907×10^{-6} and 5.732×10^{-8}, respectively, in euros in auctions held in May 1998 – December 2000. I convert these at the average EUR/GBP exchange rate in the period of 1.55 (where EUR/GBP in 1998 is (1/6.55957)FRF/GBP). The average maximum supply in my study is £4.17 billion, so the estimates imply $\rho = \frac{(6.907\times10^{-6})}{1.55} \times 4.17 \times 10^9$ and $\rho = \frac{(5.732\times10^{-8})}{1.55} \times 4.17 \times 10^9$, respectively.

Boyarchenko et al. (2021) calibrate a CARA parameter of 366.61 given a total auction supply normalized to one, so $\rho = 366.61\hat{Q}$, where $\hat{Q}$ is the ratio of the average maximum supply in my model (£4.17 billion) to the average issuance in their study (£17.3 billion, converted at the exchange rate on the issuance dates of the US Treasury notes in their sample). So $\hat{Q} = 4.17/17.3 = 0.24$ and $\rho \approx 366.61 \times 0.24 = 88.0$.

Allen and Wittwer (2023) provide a median estimate for the risk aversion parameter in their setting equal to 0.006, for a total auction supply normalized to one. The auction supply is on average 4.12 billion CAD, which is approximately £2.43 billion (converted from CAD at an exchange rate of 0.59). Following the same approach as above, this implies $\hat{Q} = 4.17/2.43 = 1.72$, and therefore $\rho \approx 0.006 \times 1.72 = 0.0103$.