Essays on auctions in financial markets

Charlotte Grace
Nuffield College, University of Oxford

September, 2023



A thesis submitted for the degree of
Doctor of Philosophy in Economics

Essays on auctions in financial markets

Charlotte Grace
Nuffield College, University of Oxford

September, 2023

Abstract

This thesis comprises four chapters on the topics of bidding behaviour and welfare in auctions in financial markets. The first chapter examines the behaviour of participants in the Bank of England's liquidity auctions. It provides a method to compare the performance of alternative models of behaviour. It finds that a model of "truthful bidding"—bidding one's true value—outperforms a conventional strategic model in which bidders shade their bids to maximise their expected surpluses. The second chapter provides a complementary approach to analyse bidding behaviour in the liquidity auctions. The approach is motivated by the fact that bidders submit few bids per auction, at odds with the predictions of the conventional model. It estimates the surplus that bidders forgo by submitting so few bids. The results suggest that the conventional strategic model does not characterise this aspect of bidding behaviour well. The third chapter compares the performance of the mechanism used in the Bank of England's liquidity auctions to simpler alternatives. It shows that the design increases welfare, and that this welfare gain can be large relative to the comparators. The results show that the Bank of England always benefitted from the design but the effect on the bidders was ambiguous. The final chapter provides conditions under which increasing efficiency by combining competitive auctions for substitute goods benefits both the bidders and the auctioneer. It also provides conditions under which each side of the market loses. It estimates the magnitude of these effects in Mexico's domestic government bond market.[1]

[1]This thesis contains approximately 75,000 words.

Declaration

This thesis comprises four chapters. Chapters 1, 2 and 4 are solely the product of my own work. Chapter 3 is jointly authored with Dr Julia Giese (Bank of England). I conducted all of the empirical analysis in the chapter.

The views expressed in this thesis are solely mine and should not be interpreted as representing the views of the Bank of England, Banco de México or Gobierno de México.

Charlotte Grace (*September 2023*)

Contents

Introduction **1**

Chapter 1 Explaining bidding in the Bank of England's liquidity auctions **7**

1 Introduction **8**

2 Institutional setting and data **12**
2.1 Indexed Long-Term Repo auction 12
2.2 Data 15
2.3 Bidding behaviour by bidder size 18
2.4 Bids submitted per bidder 19

3 Candidate models of behaviour **21**
3.1 The economic environment 21
3.2 Bid shading in equilibrium 27

4 Econometric method **33**
4.1 Identification 33
4.2 Estimation 36
4.3 Testing procedure 38

5 Results **43**
5.1 Model fit 43
5.2 Rationalising truthful bidding 45

6 Conclusion **50**

Appendices **52**
A Proof of Proposition 1 52
B Robustness 63

Chapter 2 An alternative approach to evaluate bidding in the Bank of England's liquidity auctions **66**

1 Introduction **67**

2 Interpretation of bidding costs **69**

3 Set identification of the marginal valuation function **70**

3.1 Price-taking bidder, no rationing 71
3.2 Price-taking bidder, allowing for rationing 72
3.3 Equilibrium conditions 76

4 Estimating marginal valuation functions **78**

5 Estimating bidding costs **80**

6 Results **81**

7 Conclusion **86**

Appendices **87**
A Proof of Proposition 1 87
B Proof of Corollary 1 94

Chapter 3 Comparing the product-mix auction to alternatives **96**

1 Introduction **97**
1.1 Welfare results 100
1.2 Distribution of surplus 101

2 Background **103**

3 Method **108**
3.1 Truthful bidding 108
3.2 Alternative auction rules 109
3.3 Surpluses, and welfare 111

4 The SSA **112**
4.1 Separate auctions run sequentially 117

5 The RPA **118**

6 Welfare and surpluses in periods of stress **124**

7 Alternative objective functions **125**

8 Broader benefits of the PMA **126**

9 Conclusion **127**

Appendices **129**
A The Bank of England's algorithms 129
B Additional figures 135

Chapter 4 Who gains from efficient auctions? **143**

1 Introduction **144**

2 Model **148**

3 Bidder surplus **151**
3.1 Bidder Uncertainty Effect 152
3.2 Bidder Allocation Effect 156
3.3 Market-specific bidder surplus 157

4 Auctioneer surplus **157**
4.1 Auctioneer Price Level Effect 161
4.2 Auctioneer Relative Price Effect 162

5 Special cases **163**
5.1 The symmetric case 164
5.2 Standard functional forms 165
5.3 Constant relative marginal costs 168

6 Extensions **168**
6.1 Profit-maximising auctions 168
6.2 Variable total supply 170
6.3 Alternative mechanism: reference price auction 171

7 Primary market for Mexican government debt **173**
7.1 Institutional details and data 173
7.2 Method 176
7.3 Results 180

8 Conclusion **185**

Appendices **187**
A Proofs 187
B Labelling the goods 195
C Distribution of bidder surplus across markets 195

D Decomposition of auctioneer surplus 196
E Example of Proposition 4 199
F Generalised Pareto distribution 200
G Mexican government debt auctions: alternative estimation method 202
H Spanish renewable energy procurement: model predictions 205

Conclusion **207**

References **208**

Introduction

Overview This thesis comprises four chapters on the topics of bidding behaviour and welfare in auctions in financial markets. The first two chapters examine the behaviour of participants in the Bank of England's (BoE) liquidity auctions and analyse whether bidders act strategically. The third chapter evaluates the distributional effects of the BoE's novel auction design, which jointly allocates multiple goods. The fourth chapter analyses the distributional effect of combining auctions from a theoretical perspective. It also estimates the magnitude of these effects in Mexico's domestic government bond market.

Chapter 1 The first chapter investigates bidding behaviour in the BoE's liquidity auctions. These auctions provide liquidity insurance to the financial system by allocating funds to financial institutions, lent against various collateral assets. For each type of collateral, a divisible quantity of funds is allocated at a uniform price. If participants bid their true values for loans, i.e. "truthfully", the design, a product-mix auction (PMA), maximises welfare by allowing the quantity of funds lent against different types of collateral to react flexibly to market conditions. Moreover, if bids do correspond to values, the design directly provides information about demand for liquidity and market stress.

However, it is not clear ex ante that participants will bid truthfully. In a conventional model in which bidders are strategic and risk neutral, they may maximise their surpluses by submitting bids that differ from their values. In particular, bidders with larger demands have an incentive to submit bids below their true values. This is because a bidder's bid for an additional unit may affect the auction price, which they pay for all units that they win.

This chapter provides a method to test the relative performance of alternative models of behaviour, which uses data on participants' bids but does not require data on their values. The testing procedure compares the pairwise relative fit of candidate models. I apply Backus, Conlon and Sinkinson's (2021) method, which combines Rivers and Vuong's (2002) comparison of non-nested alternatives and Berry and Haile's (2014) exclusion restrictions. The test is based on the correlation between the values implied by a model and variables that affect the optimal bid shading but are uncorrelated with bidders' true values. The procedure has not yet been adapted to an auction setting.

I compare a conventional model of bidding, developed by Kastl (2011), in which bidders shade their bids to maximise their expected surplus, to a "truthful bidding" model in which

bidders bid their true values. The results show that behaviour in the BoE's liquidity auctions is better explained by truthful bidding than by the conventional strategic model.

I give two explanations for this result: risk aversion and the costs of determining the optimal strategy.

First, I generalise the conventional model to allow for risk aversion by deriving novel identifying conditions which recover bidders' values under an assumed degree of constant absolute risk aversion. 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.) The results show that truthful bidding can be rationalised within the conventional model by a degree of risk aversion which is consistent with that found in studies in similar contexts (Armantier and Sbaï, 2006; Boyarchenko, Lucca and Veldkamp, 2021).

Second, truthful bidding may be explained by the fact that the optimal strategy can be complex and the benefits to the participants of bidding strategically can be small. Truthful bidding is a best response when the cost of calculating the otherwise optimal strategy outweighs the additional expected surplus generated by that strategy. My estimate of a lower bound on the size of these costs is the difference between the surplus from the otherwise 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. This implies that, if bidders' actual strategies are to bid truthfully, they obtain up to 95% of the surplus that would have been generated by an approximate best response.

Overall, the chapter suggests that a simple model of truthful bidding might be a reasonable approximation of behaviour in the Bank of England's liquidity auctions.

Chapter 2 The second chapter provides an alternative approach to analyse the conventional strategic model of bidding behaviour in the Bank of England's liquidity auctions. The model, described in Kastl (2011), assumes that bidders face costs of submitting additional bids, or 'bidding costs'. These costs aim to explain the fact that bidders often submit very few bids per auction, which is difficult to explain in the absence of these frictions because we would expect bidders to fine-tune their individual demand curve (by submitting multiple bids) in order to maximise their expected surpluses.

The measure of model fit that I focus on is the size of the bidding costs required to explain the data under the assumptions of the model. To estimate the bidding costs, I make use of the fact that, given the bidding costs, the number of bids to submit is a strategic choice for bidders, alongside the quantity and price of each bid. To determine the optimal number of bids, a bidder trades 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, submitting just a few bids is optimal.

Estimating the bidding costs requires estimating bidders' full marginal valuation functions. In contrast, the testing procedure described in Chapter 1 only requires estimates of bidders' marginal values at the quantities they demand at each bid price. I therefore derive novel necessary conditions for bidders' values in the conventional model of bidding, which set-identify bidders' marginal valuation functions. I combine these with the necessary conditions described in Chapter 1 to estimate the valuation functions implied by the model. These are used to estimate the expected marginal benefit of submitting an additional bid, which provides a lower bound on the bidding costs implied by the model.

I find that the estimated bidding costs for bidders are relatively large because bidders submit so few bids per auction. However, the physical costs of submitting bids are trivial in the liquidity auctions. This suggests that, under the interpretation of bidding costs as physical costs, the model does not fit this dimension of behaviour well.

Chapter 3 The third chapter compares the Product-Mix Auction (PMA), the mechanism used by the Bank of England for its liquidity auctions, to simpler alternative auction designs. The PMA, described in Klemperer (2010, 2018), allows for the quantity of funds lent against different types of collateral to react flexibly to the economic environment and market stress. To evaluate the benefits of this flexibility, I compare the design to (i) separate simultaneous auctions (SSA) for loans secured by the different sets of collateral, each with a fixed supply of funds and (ii) a "reference price auction" (RPA) in which the auctioneer fixes a pair of notional prices for the two collateral sets, i.e., "reference prices", and the highest bids relative to their reference prices are accepted. These two comparators are simpler, reasonable alternatives, used both historically and by other central banks.

The PMA jointly determines the quantities of funds lent against the two types of collateral assets. It finds the competitive equilibrium, assuming both that the Bank of England's supply curves express the prices it is willing to accept to supply the loans and that bids express bidders' true valuations for loans. Under these assumptions, the PMA maximises

welfare because it finds the competitive equilibrium and implements this by choosing an allocation such that marginal welfare is the same for both collateral sets. The alternative auction designs may fail to maximise welfare because of the constraints they introduce: the SSA and RPA do not permit quantities and the price difference, respectively, to adjust to the participants' or Bank of England's preferences.

The main finding is that the PMA increased welfare (the difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept) by approximately 50%, or 2 basis points per loan, relative to these two alternative designs in the sample period of June 2010 – January 2014. The RPA could not have maximised welfare with any fixed reference prices because of variation in the optimal price difference over the sample. And although one particular SSA could have almost maximised welfare in the period, this result was specific to the stability of the optimal quantities in the particular sample period. Moreover, selecting this best-performing SSA would have required the Bank of England to accurately forecast the optimal quantities to supply against each collateral set, which seems unrealistic, and the welfare losses of other SSAs were large.

A second finding is that using the PMA always gave the BoE more or at least as much surplus and revenue as running any possible SSA or RPA. The effect on the bidders, in aggregate, was ambiguous. Bidders, on average, shared in the gains only when the fixed share supplied against the less liquid collateral used in the SSA was low, or when the difference in fixed reference prices between the two collateral used in the RPA was high, so that the PMA typically lent more against the less liquid collateral relative to the alternative.

Three features in the period can explain why the bidders did not always gain. First, there was no significant stress. Second, demand for only one set of collateral was not flat. Third, there was little use of the "paired bidding" option in the design. Absent any one of these three features, the bidders, as well as the BoE, might have benefited more from the use of the PMA.

I expect, and confirm in simulations, that the welfare gain and the benefits to the BoE of the PMA would be qualitatively similar, but quantitatively larger in absolute terms, in a less stable period than the period studied.

Importantly, neither the SSA nor the RPA would have achieved the BoE's main objectives in setting up the PMA. The auctions were intended to improve the availability and flexibility of liquidity insurance to the banking system as a whole, while mitigating potential moral hazard, and informing the BoE about stress in the market (Fisher, 2011). These broader

benefits of the auctions are taken into account, to the extent that the BoE's supply curve internalises some externalities. The automatic adjustment to market conditions of both the amount of funds loaned in the PMA, and the interest rate premium charged, improves efficiency not only as measured in this chapter, but also in this broader sense, too.

Chapter 4 The final chapter analyses the impact of combining auctions into a single optimal mechanism on the distribution of surplus between bidders and the auctioneer. In my setting, participants bid competitively for substitute goods and the optimal mechanism is a Product-Mix Auction (PMA). I provide conditions under which the optimal mechanism increases both expected bidder and auctioneer surplus or reduces one of them. I then illustrate the model's predictions using novel data from Mexico's domestic government bond market.

Using one auction to jointly allocate multiple substitute goods is increasingly popular among public sector institutions, especially in energy markets and financial markets. Running a multi-object auction, in which participants bid competitively for substitutable goods, increases efficiency, but the effect on the distribution of surplus between bidders and the auctioneer is less clear. This chapter analyses this question both theoretically and empirically.

I describe a model in which an auctioneer wishes to allocate a fixed supply across two goods that she views as imperfect substitutes. The objective is to maximise welfare. For each good, there is a large number of risk-neutral bidders, each with unit demand, who therefore bid their true values. The optimal mechanism is therefore a PMA. I compare the effects on bidder and auctioneer surplus of running this optimal mechanism relative to running an independent auction for each good with a fixed supply.

The results show that the optimal mechanism increases efficiency, yet it does not necessarily benefit both sides of the market. I provide conditions under which the optimal mechanism reduces expected bidder surplus, and conditions under which it reduces expected auctioneer surplus.

Nonetheless, the conditions required for either side to lose are fairly extreme. In the symmetric case of my model (in which the bidders' aggregate preferences, and the auctioneer's preferences, are symmetric in the two goods), both bidders and the auctioneer benefit from the optimal mechanism if the demand curves are convex but log concave. This is because bidders as a group gain from the uncertainty in the quantities sold in the optimal mechanism, but the uncertainty also translates into a higher expected payment to the auctioneer. If the

demand curves are log convex, only the auctioneer benefits from the optimal mechanism. Alternatively, if the demand curves are very concave, only the bidders benefit. I also show how these results extend to the asymmetric case.

I also estimate the magnitude of these distributional effects in the primary market for Mexican Treasury bills. Using bid-level data from the auctions for Mexican Federal Treasury Certificates (CETES), which are held independently for each bond, I estimate the difference in bidder and auctioneer surpluses between running a counterfactual PMA and the existing design. I find that combining the auctions of 28-day and 91-day bonds, and combining the auctions of 182-day and 364-day bonds, both increase both bidder and auctioneer surplus. This analysis provides support for the view that the optimal mechanism typically benefits both sides of the market.

The favourable distributional effects that I find, both in theory and in practice, may give additional reason for combining competitive auctions of substitute goods, beyond the overall efficiency gain.

CHAPTER 1
Explaining bidding in the Bank of England's liquidity auctions

Abstract

This chapter provides a method for comparing the performance of different models of bidding behaviour. It uses data on participants' bids but does not require data on their values. In the Bank of England's uniform-price divisible-good liquidity auctions, I find that a model of "truthful bidding"—bidding one's true value for liquidity—outperforms a conventional model in which bidders shade their bids to maximise their expected surpluses. I provide two possible explanations for this result. First, when bidders are sufficiently risk averse, optimal strategies in the conventional model approximate truthful bidding. For the conventional model, I develop new identifying conditions which allow for risk aversion. I find that the degree of risk aversion required for truthful bidding to be approximately optimal is consistent with that found in studies in similar contexts. Second, the optimal strategy can be complicated. Truthful bidding is preferable if the cost of calculating what would otherwise be the optimal strategy exceeds around 5% of bidder surplus.

1 Introduction

Optimal bidding strategies in divisible-good auctions can be both risky and complex. A participant's bid expresses the price she is willing to pay as a function of the quantity she demands. In uniform-price auctions—the focus of this chapter—a participant's bid for a marginal unit may affect the price she pays for inframarginal ones.[1] This creates an incentive for the participant to bid differently from her true marginal value for the good at each unit, and the magnitude of this optimal "bid shading" varies across the quantity that she demands. It depends on the participant's risk preferences—bidding below her value reduces the expected payment that she will make if she wins, at the cost of reducing the probability of winning at all—as well as on the participant's beliefs over the economic environment and on her constraints.

Bid shading in practice is difficult to infer: participants' true values are rarely observed. Empirical studies of divisible-good auctions therefore typically assume a model of bidding when assessing efficiency, revenue and alternative auction designs (see, e.g., surveys by Gentry et al., 2018; Hortaçsu and McAdams, 2018; Hortaçsu and Perrigne, 2021). The nature of bidding behaviour affects these market outcomes so the choice of model is critical to the analysis.

This chapter provides a method to test the relative performance of alternative models of behaviour. It uses data on participants' bids but does not require data on their values. I analyse the Bank of England's Indexed Long-Term Repo (ILTR) auctions, in which the central bank provides liquidity insurance to financial institutions by lending a divisible quantity of funds against collateral. The payment rule is uniform-pricing.

I find that bidding behaviour is 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, participants submit bid functions corresponding to their true marginal valuation functions. Conversely, in a conventional strategic model, participants best respond to their beliefs about the economic environment and other participants' behaviour. In the conventional model that is most appropriate in my context, participants may optimally bid above or below their marginal values for liquidity, so truthful bidding is a natural comparator.

[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 her bid for each unit that she wins. Her bid therefore also affects both her probability of winning and the expected payment.

Truthful bidding can be explained 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.

The relative fit of alternative models of bidding behaviour is tested using exclusion restrictions which exploit variation in the amount of bid shading implied by the different models. Variables that affect bidders' bid shading but not their values for liquidity are selected as instruments. For example, with private values, the number of other bidders in the auction affects one's market power but not one's valuation. These instruments provide variation in the amount of bid shading implied by the correct model which is independent of bidders' values for the good. Conversely, under a misspecified model, correlation between the instruments and model-implied bid shading will induce correlation between the instruments and model-implied values. Stronger correlation between the instruments and model-implied values is therefore evidence of worse model fit.

The testing procedure to compare the pairwise relative fit of candidate models is based on these correlations. It asks only which model is *more* consistent with the data and is robust to misspecification of the models. The procedure follows Backus, Conlon and Sinkinson's (2021) analysis of firm conduct, which extends Rivers and Vuong's (2002) comparison of non-nested alternatives using the exclusion restrictions proposed by Berry and Haile (2014). While I focus on uniform-price divisible-good auctions, this 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 behaviour, 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, Hortaçsu and Kastl, 2013; Hortaçsu, Kastl and Zhang, 2018; Allen, Kastl and Wittwer, 2022). This is the most appropriate benchmark to analyse behaviour in the Bank of England'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.

I consider two explanations for why a model of truthful bidding outperforms this conventional model of behaviour, 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 Bank of England's liquidity auctions. First, 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). Second, 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] Unsurprisingly, when bidders are risk averse, optimal strategies in the conventional model are closer to truthful bidding—if a bidder is sufficiently risk averse, she does not want to gamble on losing at prices at which she would prefer to win (or winning at prices at which she would prefer to lose) by submitting a bid that differs from her 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 rationalised within the conventional model by a degree of risk aversion which is consistent with that found in studies in similar contexts (Armantier and Sbaï, 2006; Boyarchenko, Lucca and Veldkamp, 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% of the surplus that would have been generated by an approximate

[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.

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 chapter primarily contributes to the literature on testing equilibrium behaviour 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.[3] By contrast, I exploit exclusion restrictions to compare the goodness-of-fit of alternative models of behaviour. I therefore do not require over-identifying moments, but 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 is also robust to model misspecification. It has been used to test firm conduct in product markets (Backus, Conlon and Sinkinson, 2021; Duarte, Magnolfi and Roncoroni, 2021; Starc and Wollmann, 2022) and labour markets (Roussille and Scuderi, 2022). To my knowledge, it has not yet been adapted to an auction setting.[4]

Other studies estimate marginal values using data outside the model and compare these directly to the observed bids. For example, Wolfram (1999) and Borenstein, Bushnell and Wolak (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 maximisation, yet smaller bidders persistently deviate from equilibrium bidding. Similar to this chapter, Hortaçsu and Puller (2008) show that this can be rationalised by a cost of sophistication that only larger bidders are willing to incur.[5] In contrast to these studies, my testing procedure does not require data on marginal values, which are often unavailable.

This chapter also contributes to the literature on risk aversion in financial markets. Armantier and Sbaï (2006) and Boyarchenko, Lucca and Veldkamp (2021) estimate the degree

[3]Wolak (2007) fails to reject profit maximisation by participants of Australian electricity markets, while Chapman, McAdams and Paarsch (2006) apply McAdams' (2008) method to the Bank of Canada's auctions of cash reserves and find that bids were "close to" best response (i.e. an ε-equilibrium).

[4]Roussille and Scuderi (2022) model their labour 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.

[5]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. The parameters of the cognitive hierarchy model are unidentified given only bidding data, so I cannot assess it with my approach. Truthful bidding is of course a special case in which Level-0 behaviour is defined by truthful bidding and adopted by all bidders.

of risk aversion of bidders in French and US Treasury auctions, respectively. Applying their findings to the Bank of England's liquidity auctions would suggest that truthful bidding is approximately optimal in our context. Using data from Indian Treasury bill auctions, Gupta and Lamba (2022) use risk aversion to explain the combination of increased participation and higher uncertainty during the 2013 "taper tantrum". They build on Kastl's (2011) model by assuming bidders have constant relative risk aversion (CRRA), and identify values by assuming both that bidders' marginal valuation functions are step functions and that at most one bidder submits a bid equal to each possible auction price (i.e. bidders do not tie on the margin). I instead permit bidders to have constant absolute risk aversion (CARA) and allow for ties on the margin, which are highly probable in my context and in others in which the set of possible bid prices is discrete.[6]

The chapter proceeds as follows. Section 2 provides background for the Bank of England's liquidity auctions and introduces the data. In Section 3, I describe the alternative models of behaviour and their implications for the amount that bids differ from bidders' values. The identification of bidders' values, the estimation method and the testing procedure are described in Section 4. Section 5 gives the results of the pairwise comparisons of models and shows how truthful bidding can be explained. Section 6 concludes.

2 Institutional setting and data

2.1 Indexed Long-Term Repo auction

My empirical setting is the Bank of England's (BoE) Indexed Long-Term Repo (ILTR) auction. This 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 range of BoE facilities that provide liquidity insurance. I study the period

[6]CARA implies that a bidder's risk aversion does not vary with her net surplus from the auction (defined by the area between her marginal valuation function and the auction price, up to the quantity she is allocated), whereas Gupta and Lamba's (2022) CRRA specification implies that the bidder's risk aversion is decreasing in her net surplus from the auction. CARA seems more appropriate in my context for at least two reasons. First, the financial institution's net surplus from the auction is trivial relative to its total assets so CARA seems to be a more suitable 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, it is not clear why their risk aversion would be lower when the stakes were higher. Moreover, if the manager has a concave utility function because of the remuneration structure, CRRA would imply that the manager is rewarded more for winning in the auction when the participant's net surplus from the auction is smaller. It is unclear whether this would be an appropriate assumption.

June 2010 – January 2014, in which the operations were held monthly.[7]

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 collaterised. 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.[8] 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 (PMA), originally developed in Klemperer (2008) 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.

Bidders 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.[9] Nixon (2014) shows that the number of institutions with OMO access was fairly stable from October 2009, before the ILTRs were introduced, to January 2014, ranging from 48 to 52. There was no formal obligation for registered bidders to participate in the auctions.

Bids Bidders can submit any number of sealed 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

[7]The auction design, supply curves representing the BoE's preferences between collateral sets, and eligibility for entry have evolved over time. These auctions are currently held weekly, with rules that have been modified since the sample period. Further details can be found at `https://www.bankofengland.co.uk/markets/bank-of-england-market-operations-guide/our-tools`.

[8]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`.

[9]Since 2014, after the sample period, broker-dealers and central counterparties are eligible to register for the SMF (de Roure and McLaren, 2021).

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 (Bank of England, 2010a).

Supply Prior to the auction, the BoE commits to two supply curves:

1. The maximum supply made available across goods A and B is publicly announced prior to the auction.

2. The BoE commits to a privately known "relative supply" curve. It is measured by the difference between the prices of goods B and A and is an increasing function of the quantity allocated of good B. This is "pinned down by [the BoE's] preferences" to provide liquidity insurance at prices which incentivise prudent liquidity management (Fisher, 2011a, p.12).

Market clearing The PMA 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 sealed bids and BoE's supply preferences, and so is uncertain from the perspective of bidders.

For each good, the auction price is the highest losing bid for that good.[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.

Information Certain information about the auction results is made public: the total

[10] In the sample period, 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. She is allocated the good which maximises her surplus, given the clearing prices and assuming bids reflect bidders' true preferences. For example, a paired bid may specify demand of £50 million at a spread of 2bps for good A or 25bps for good B. If the clearing prices are 1bps and 22bps respectively, the bidder is allocated £50 million of good B. In practice, bidders rarely made use of paired bids so I drop them from the sample.

[11] The precise price determination rules (where competitive equilibrium does not determine the prices uniquely) change slightly over the period, from lowest-winner pricing (LWP) to highest-loser pricing (HLP). Under LWP, the auction price for each good is equal to the lowest winning bid on that good. Under HLP, the auction price for good A is equal to the highest losing bid on that good, or 0bps if all bids are fully allocated. The auction price for good B under HLP is equal to the maximum of the highest losing bid on good B and the minimum price that the BoE is willing to accept for the quantity allocated of good B, determined by the auction price for good A and the relative supply curve.

In the majority of cases, this change is irrelevant, as the marginal bid is rationed so that the two prices are identical.

quantity of bids for each good, total amounts allocated, clearing spreads and rationing coefficients. None of the number of bidders, the number of individual bids, and the range of bid prices are revealed.

2.2 Data

I use a unique proprietary dataset that consists of all bids submitted in the BoE's ILTR auctions, held monthly in June 2010 – January 2014, and the relative supply curves used by the BoE to determine the allocation across goods. To my knowledge, this data has only once been analysed.[12]

Supply The maximum supply, publicly announced prior to the auction, was £5 billion for the 3-month term auction and £2.5 billion for the 6-month term auctions. The minimum auction prices of goods A and B were 0bps and 5bps, except for the prices in the 6-month auctions from May 2011, in which the minimum for good B 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 cannot disclose their precise functional forms.

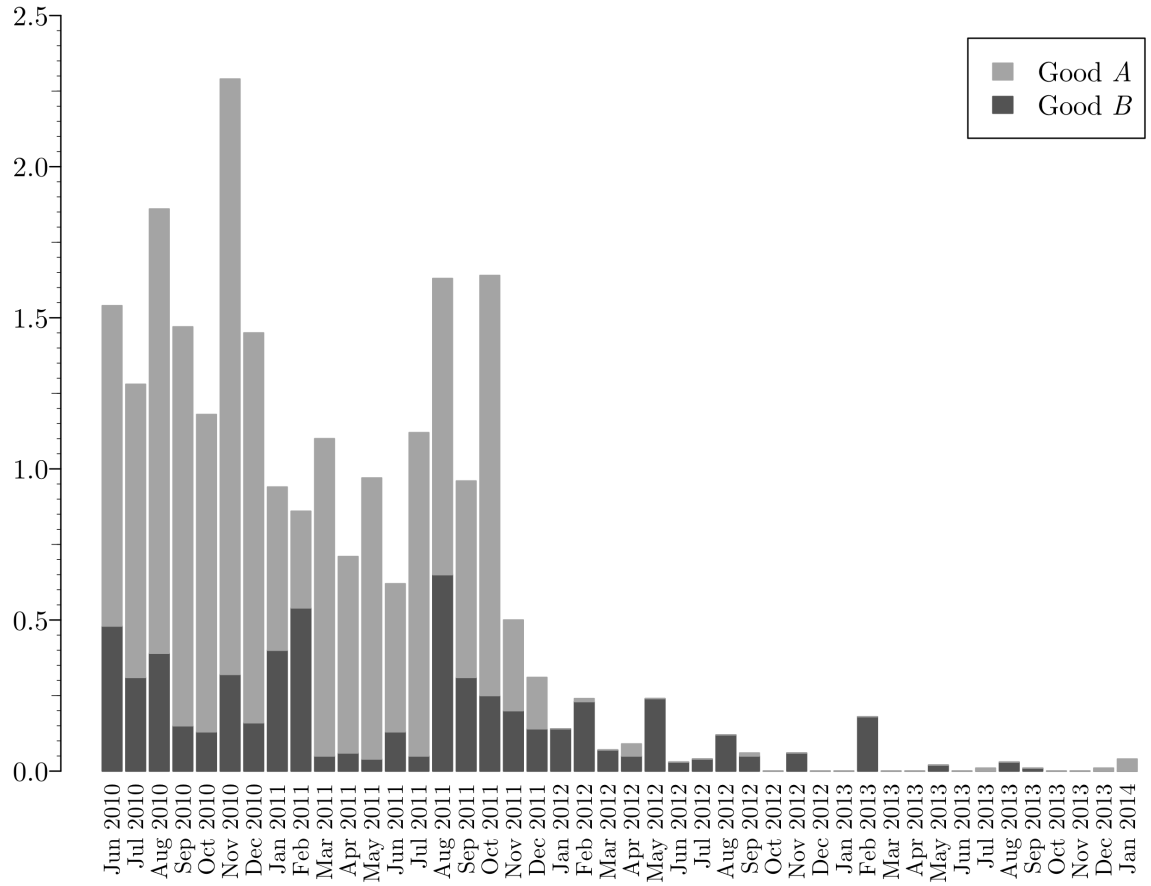
Bids The bidding data is de-identified. Across auctions, unique participants are indexed from 1 to P. For each participant, I observe the full set of bids that the individual submits on goods A and B, if any, in each of the 44 auctions. Each bid consists of the type of collateral the bidder will provide (the good), the spread in basis points over Bank Rate they are willing to pay (the bid price) and the quantity of liquidity they demand in £million. For each good, a bidder's bids imply a (price, quantity)-schedule, i.e. an individual bid function, which is a step function.

Aggregate demand The total quantity of funds demanded in the auctions decreased substantially in the second half of the sample period (see Figure 1) so I limit the analysis to the first 24 months of auctions (from June 2010 to May 2012).[13] Aggregate demand also varies over the period, which the estimation procedure described in Section 4.2 aims to account for.

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

[13] Winters (2012) provides three explanations for this. First, an increase in the availability of funding from the BoE's quantitative easing programme and cheaper funding in the market than from the BoE reduced demand for good A. Second, a reduction in participants' holdings of assets in the Level B collateral set reduced demand for good B. And finally, funding allocated in the ILTR auctions was of relatively short maturity, so was not well-suited to meet participants' regulatory liquidity requirements. Increased availability of funds from the euro area and from alternative facilities in the BoE also likely played a role. For example, the Funding for Lending Scheme was introduced by the BoE in July 2012, allowing borrowing of up to four years. Another source is the BoE's Extended Collateral Term Repo (ECTR) operation, consisting of monthly auctions for 6-month term liquidity, although this was only activated once in June 2012.

Figure 1: Aggregate quantity demanded as a share of the maximum supply, June 2010–January 2014.



Summary statistics Table 1 provides summary statistics of the data. The auction price of good A is typically 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.

Table 1: Summary statistics for ILTR auctions in June 2010 – May 2012

	3-month term		6-month term	
Number of auctions	16		8	
Both goods				
	Mean	Std. Dev.	Mean	Std. Dev.
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
Good *A*				
	Mean	Std. Dev.	Mean	Std. Dev.
Quantity of funds demanded (% of allocation of good A)	125.72	22.68	171.98	66.62
Bid price (weighted, basis points)	2.51	2.85	3.06	4.31
Number of bids submitted per bidder (unweighted)	1.52	0.84	1.49	0.86
Number of bids submitted per bidder (weighted)	1.79	0.94	1.86	1.04
Quantity demanded by bidder (% of maximum supply)	13.65	11.87	13.93	12.30
Quantity of funds allocated (% of total allocated)	83.72	12.00	70.65	20.19
Price paid (weighted, basis points)	1.14	1.71	0.59	0.79
Good *B*				
	Mean	Std. Dev.	Mean	Std. Dev.
Quantity of funds demanded (% of allocation of good B)	151.26	58.15	202.74	91.49
Bid price (weighted, basis points)	23.61	9.96	38.96	16.15
Number of bids submitted per bidder (unweighted)	1.64	1.31	2.30	1.71
Number of bids submitted per bidder (weighted)	2.90	2.11	2.95	2.28
Quantity demanded by bidder (% of maximum supply)	4.30	6.41	9.69	9.87
Quantity of funds allocated (% of total allocated)	16.28	12.00	29.35	20.19
Price paid (weighted, basis points)	21.67	6.52	43.43	1.37

Weightings: 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; and the number of bids submitted per bidder is weighted by the total quantities demanded by bidders.

Table 2: Summary statistics by bidder size in the ILTR auctions in June 2010 – May 2012

Good *A*								
	3-month term				6-month term			
	Small bidders		Large bidders		Small bidders		Large bidders	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Bid price (weighted, basis points)	3.22	2.13	2.51	2.85	5.33	5.20	3.03	4.29
Number of bids submitted per bidder (unweighted)	1.09	0.29	1.67	0.91	1.25	0.45	1.58	0.96
Good *B*								
	3-month term				6-month term			
	Small bidders		Large bidders		Small bidders		Large bidders	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Bid price (weighted, basis points)	24.42	13.93	23.47	9.10	41.36	23.49	38.88	15.85
Number of bids submitted per bidder (unweighted)	1.28	0.72	2.14	1.74	1.89	1.17	2.50	1.92

Definitions: 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.

2.3 Bidding behaviour by bidder size

Bidders differ in the amount that they demand, and therefore in their ability to affect the auction price. In conventional models of bidding in uniform-price auctions, larger bidders have incentives which correspond to the quantity-shading incentives of oligopolists facing uncertain demand in Klemperer and Meyer (1989). A bidder can reduce the expected auction price paid for all units that she wins by reducing the quantity that she demands at each price; the strength of this incentive is increasing in the quantity that she demands. 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 2 shows the average bid prices of bidders who demand less than 2.5% of the maximum supply in all auctions, who are labelled "Small", and bidders in the remaining set, labelled "Large".[14] Larger bidders do submit lower bid prices in the ILTR auctions, but this does not isolate differences in strategic incentives from unobserved variation in bidders' 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 chapter is to employ a testing procedure, described in Section 4.3, which isolates the variation in bidders' strategic incentives from variation in their values in order to evaluate alternative models of behaviour.

[14]Some natural alternative definitions of bidder size give similar results.

2.4 Bids submitted per bidder

A striking feature of the data is that bidders submit few bids per auction. Figure 2 plots the density of bids per bidder per auction for goods A and B, unweighted and weighted by the quantities demanded by bidders.

This characteristic of bidding is also seen in many other multi-unit auction settings.[15] It motivates the model of bidding behaviour, developed by Kastl (2011, 2012), which is now widely used in empirical studies (e.g. Cassola, Hortaçsu and Kastl, 2013; Hortaçsu, Kastl and Zhang, 2018; Allen, Kastl and Wittwer, 2022).

In the model, bidders face costs of submitting additional bids, which I refer to as "bidding costs". With these frictions, the number of bids to submit is also a strategic choice: a bidder must trade off the marginal bidding cost with the marginal benefit of fine-tuning her 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 2), 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 good, we would expect her to submit a large number of bids, as her incentive to strategically understate her value is increasing in the quantity that she demands (see Section 2.3).

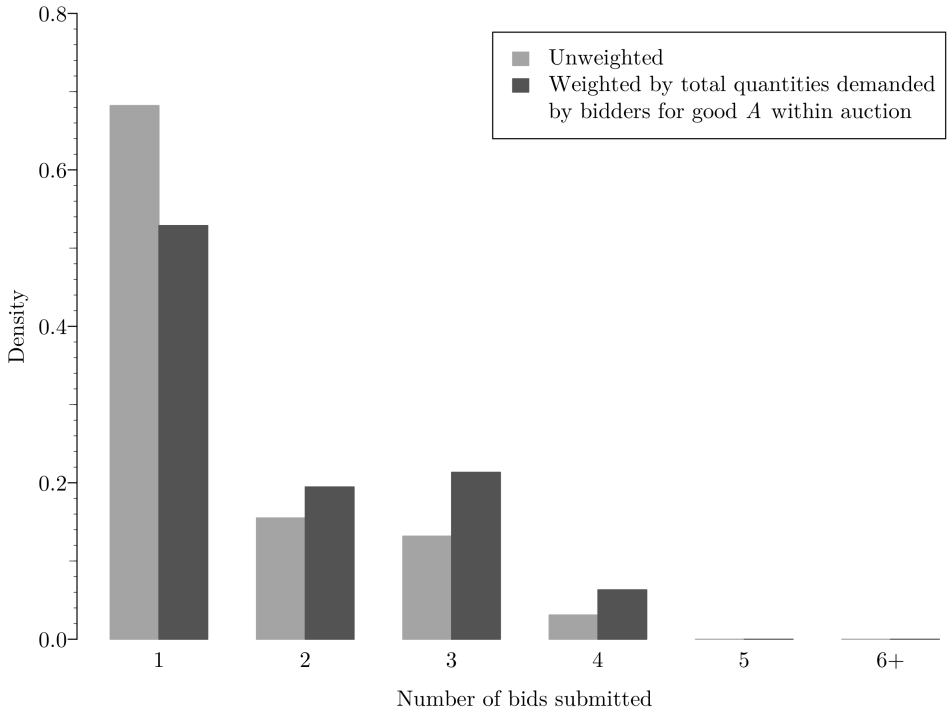
It is therefore natural to use Kastl's (2011) model with bidding costs as a benchmark to compare alternative models of behaviour.

[15] 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, Kastl and Zhang, 2018), respectively. In wholesale electricity auctions, Hortaçsu and Puller (2008, p.106) document that bidders in Texas "do not make full use of the strategy space available to them, but rather use coarse-grained bidding strategies", and bidders only submit 4.4 bids on average in Spain (Reguant, 2014). This small number of bids is also observed in central bank operations: Cassola, Hortaçsu and Kastl (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, McAdams and Paarsch, 2006).

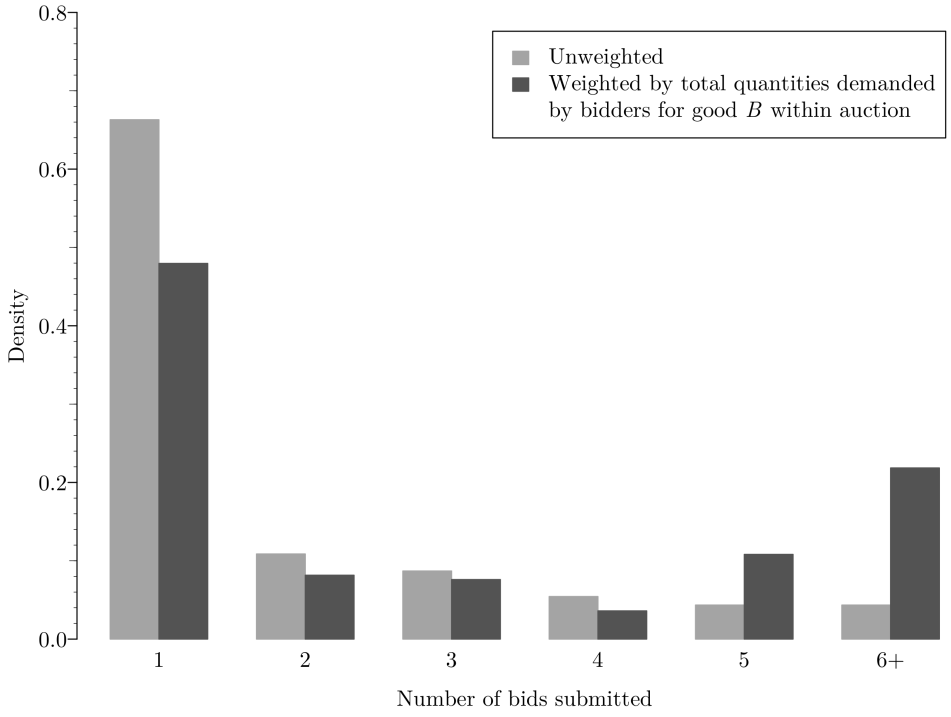
Moreover, institutional constraints on the number of bids that a bidder may submit are almost never binding. For example, both Kastl (2011) and Hortaçsu and Puller (2008) document that not one bidder submits the maximum number of permitted bids (10 and 40, respectively). Even when this bound is relatively tight, as is the case in the cash-management auctions studied by Chapman et al.'s (2006) in which only 4 bids are permitted, only 2.5% of bidders submit the maximum number.

Figure 2: Number of bids submitted per bidder within auction on each good in June 2010 – May 2012.

(a) Good A



(b) Good B



3 Candidate models of behaviour

My aim is to compare alternative models of bidding behaviour. Specifically, I compare a conventional model of bidding—widely used in empirical studies of multi-unit auctions—to a "truthful bidding" model in which all bidders submit bids according to their true marginal values. Moreover, I generalise the conventional model to allow for constant absolute risk aversion and show how optimal strategies in the conventional model tend towards truthful bidding as the degree of risk aversion becomes large.

In this section, I describe these models. I first detail the assumptions on the underlying economic environment which are common to the models that I consider. These closely follow the existing literature based on Kastl's (2011) framework, which modifies Wilson's (1979) share auction model to account for the fact that bidders submit a discrete number of bids. The main exception is that I relax Kastl's (2011) assumption of risk neutrality.

I then describe the additional assumptions which are specific to each model. Bidders face costs of submitting additional bids in the conventional model, which are introduced to rationalise the fact that bidders submit few bids. I model truthful bidding by simply assuming that bidders incur arbitrarily large costs if their bids differ from their marginal values.

Equipped with this setup, I show how bidders' marginal values are identified in each model and what this implies about their strategic behaviour. The following description refers to one auction, and, where obvious, I suppress the time subscript.

3.1 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 Klemperer (2008) and Section 2.1.

3.1.1 Marginal values

There are N potential bidders. Prior to the auction, each bidder $i \in \mathcal{N} = \{1, ..., N\}$, receives a privately known signal, Θ_i, which has dimension M. Let $\boldsymbol{\Theta}_{-\boldsymbol{i}} = \{\Theta_j\}_{\{j \neq i\}}$ denote the signals received by all bidders except i. The joint distribution of signals is $F(\boldsymbol{\Theta})$, where $\boldsymbol{\Theta} = \{\Theta_j\}_{\{j \in \mathcal{N}\}}$.

Bidder i's signal is distributed such that her value is above the reserve price for at most one of the goods. This greatly improves the tractability of the model and is appropriate for the majority of bidders (85%) who only bid for one good. For the purposes of the analysis, it is also appropriate for the 24% of remaining 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, so that it is irrelevant for the implications of the model.[16,17]

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 $\sum_{g=\{A,B\}} \sum_{h=\{1,2\}} N^{g,h} = N$. Bidders are assumed to be symmetric within group, that is, signals are identically distributed and marginal valuation functions are symmetric.[18]

In the model, bidders have independent private values. I cannot rule out interdependent values entirely, but three reasons suggest that this framework is a reasonable approximation.

First, while the ILTR auctions were designed to enhance financial stability by meeting demands for liquidity insurance, the sample period of 2010 – 2014 was relatively stable. This suggests that the possible interdependence, which would arise from another bidder's private signal revealing information of perceived stress in the market, was less relevant in our sample.

[16]The analysis uses local deviations in the quantity that a bidder demands to identify her values. If a bidder bids on both goods, but her 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 she demands of good B does not affect her surplus from good A (and vice versa). We can therefore treat the bidder for both goods as if she is two bidders, one for each good.

[17]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 (for details, see Klemperer, 2018), but they very rarely choose to submit such bids. The absence of these bids reveals that bidders do not view goods A and B as simple one-for-one substitutes. It is also reasonable to rule out complement and more complicated substitute preferences. This is because goods A and B represent liquidity borrowed against two different collateral sets. These can be interpreted as two goods measured in the same units that face different opportunity costs (the return on the next best use of the assets put up as collateral) but are used for broadly the same purpose of meeting liquidity needs. In this context, one-for-one or non-substitutability appear most plausible. Indeed, the BoE designed the ILTR auction, anticipating bidders to express one-for-one substitutability between goods A and B, if at all.

[18]Bidders are financial institutions with diverse liquidity needs and motivations for participation, which suggests that bidders are drawing from different distributions of signals or have different marginal valuation functions. Section 2.3 shows one example of this heterogeneity: bidders systematically differ by the quantities that they demand. In line with the literature which assumes either one or few groups (see, e.g., Hortaçsu, 2002; Kastl, 2011; Hortaçsu, Kastl and Zhang 2018; Boneva, Kastl and Zikes, 2020), I group bidders by size of demand and assume i.i.d. signals within group to permit a sufficient sample size for estimation (see Section 4.2).

Second, it is plausible that participation over the period studied was driven largely by idiosyncratic liquidity needs, rather than by speculation in the interbank market. The fairly short maturity of loans secured in the ILTR auctions limited incentives for speculation, and this was especially the case for good B, because collateral assets in the Level B set were less liquid by definition.[19] Moreover, the general collateral repo rate was a good proxy for the market value of good A, i.e. loans secured by highly liquid collateral. 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.[20]

Finally, neither Bindseil, Nyborg and Strebulaev (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.[21]

Bidders are also assumed to have utility functions which are quasi-linear in assets outside the auction. This is motivated by the fact that bidders use a negligible fraction of total assets as collateral for liquidity obtained in the ILTR auctions (de Roure and McLaren, 2021).[22]

Let $y_i \in \mathbb{R}_+$ be the quantity allocated to bidder i expressed as a share of the maximum supply. With independent private values and quasi-linear additively separable utility, bidder i's marginal valuation is only a function of y_i and her signal Θ_i. This is assumed to be non-increasing and continuous in y_i.[23]

[19]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.

[20]Tighter liquidity regulation, introduced by the UK Financial Services Authority in 2010, might also explain participation in the ILTR auctions for 6-month loans. The regulation increased demand for high-quality liquid assets with more than 3-month maturity, and required banks to regularly participate in a number of different funding markets (Banerjee and Mio, 2018). Anecdotal evidence suggests that this was one motivation for participation, but unlikely to be the key driver given the relatively short maturity of the loans (Winters, 2012).

[21]Tests for interdependent values exist for single-unit auctions, exploiting random variation in the number of bidders (Athey and Haile, 2002; Haile, Hong and Shum, 2003), but are more difficult to construct for multi-unit auctions. Bindseil, Nyborg and Strebulaev's (2009) main test exploits variation in an interbank rate which can be seen as a benchmark for the auctions. We cannot construct this in our setting given the variation in assets used as collateral, especially in the Level B set. Similarly, Hortaçsu and Kastl's (2012) test, which relies on variation specific to their institutional setting, is not applicable to this context. In their setting, some bidders are dealers, observing a subset of the other bids prior to the auction, and variation in the behaviour of these bidders can be exploited to detect interdependent values. It is not possible to observe analogous variation in my data.

[22]They find that participants on average used 0.81% of their total assets as collateral for liquidity obtained in the ILTR auctions in 2010 to 2016.

[23]The funds that a given bidder is allocated in the auction is one component of her larger stock of cash reserves. Collateral that can be used to back ILTR loans is typically deposited at the BoE prior to the

The marginal valuation of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, $v^{g,h}(y_i, \Theta_i)$, is bounded and strictly increasing in $\Theta_i \; \forall y_i$ and non-increasing and continuous in $y_i \; \forall \Theta_i$. Bidder i's signal Θ_i is distributed by the atomless function $F^{g,h}(\Theta_i)$, which has strictly positive density $f^{g,h}(\Theta_i)$ over $[0, 1]^M$.

3.1.2 Information

I now turn to a bidder's information at the time of the auction. Bids in the ILTR auctions are submitted privately through an online platform. The number of other bidders is unknown, as are their bids; bidders are assumed to have common knowledge of the number of potential bidders and the joint distribution of signals.[24]

In the ILTR setting, as in other studies in which a model of behaviour is assumed in order to recover bidders' values, the suitability of this common prior is unclear. The number of potential bidders was equal to the number of financial institutions with access to the BoE's OMOs, and bidders 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.[25]

In addition, for simplicity, I assume that the BoE's maximum supply and relative supply curve are commonly known.[26]

The number of potential bidders in each group, $N^{g,h}, g = \{A, B\}, h = \{1, 2\}$, the joint distri-

auction, although it is possible to submit collateral for settlement following the operation (Alphandary, 2014). The setup costs are therefore mostly sunk, and there are no other obvious fixed costs associated with accessing and using the liquidity. This 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. Chapman, McAdams and Paarsch (2006) corroborates this in a similar context of the Bank of Canada's liquidity auctions.

It is less clear whether the assumption of continuity is reasonable. It is supported by the fact that both the collateral used to back ILTR loans, and the funds obtained, constitute a small fraction of participants' total assets. This would be useful to explore in future work.

[24]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 on his findings.

[25]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) for further details).

[26]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) models the total supply in Turkish Treasury auctions as an AR(1) process, because bidders face some uncertainty in its precise quantity, but he finds this makes very little difference to his results. Our assumption may be more reasonable in later auctions in the sample, as some information about the relative supply function is revealed in the publicly released data of auction outcomes.

bution of signals, $F(\boldsymbol{\Theta})$, the maximum supply and the relative supply curve are commonly known.

3.1.3 Strategies

Bidders can submit up to $\bar{K}$ bids on each good, where $\bar{K}$ is arbitrarily large.[27] For each bidder, bids on each good can be summed horizontally in (price, quantity)-space to form an individual bid function, which is a step function. Hence, bidder i chooses the number of bids to submit, K_i, and, for each bid $k \in \{1, ..., K_i\}$, the bid price $b_{i,k}$, and cumulative quantity demanded, $q_{i,k}$, where the bids are ordered in increasing quantity.[28] Bid prices are non-negative integers, and quantities demanded are continuous in the unit interval.[29] 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.

A pure strategy of bidder $i \in \mathcal{N}$ is a mapping from the private signal to a bid function, $\sigma_i : \Theta_i \to \mathcal{F}_i = (K_i, \boldsymbol{b_i}, \boldsymbol{q_i})$ where

- $K_i \leq \bar{K}$.
- $\boldsymbol{b_i} \in \mathbb{R}_+^{K_i}$, $\boldsymbol{q_i} \in \mathbb{R}_+^{K_i}$, where $b_{i,k} \geq b_{i,k+1}$ and $q_{i,k} \leq q_{i,k+1}$ for $k < K_i$.
- $q_{i,k} \in [0, 1]$, $b_{i,k} \in \mathbb{N} \cup \{0\}$, and $b_{i,k} \leq \bar{K} - 1$.

3.1.4 Expected utility

I now define bidders' expected utility, which has two important features.

Firstly, I allow for bidders to be risk averse, with constant absolute risk aversion utility described by parameter $\rho \geq 0$, which is common to all bidders.

[27] In the ILTR auctions, bidders can submit any number of bids. In the model, the upper bound $\bar{K}$ is a technicality and is assumed without loss of generality as the price grid is discrete and the auction prices are always finite, so $\bar{K}$ can be set high enough such that it is never binding.

[28] And so, $(q_{i,k} - q_{i,k-1})$ is the marginal quantity demanded at bid price $b_{i,k}$.

[29] This approximates the (price, quantity)-grid in the ILTR auctions, in which the minimum permitted cumulative quantities demanded are 0.1% and 0.2% of the maximum supply in the 3-month and 6-month term auctions respectively, with increments of 0.02% and 0.04%, respectively.

I also consider two alternative assumptions on the frictions that bidders face in order to explain the facts that bidders submit bids with few steps (see Section 2.4). These assumptions differentiate the models of behaviour that I compare: a conventional "strategic bidding" model (with either risk neutrality or a specified degree of risk aversion) and a "truthful bidding" model in which bids correspond to bidders' true marginal values. I denote the friction by F^m, defined for each model m below.

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

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

where

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

The first of the two alternative assumptions on the friction F^m is that each bidder faces "bidding costs", $c_i(K_i)$, which are increasing in the number of bids that they submit. This characterises Kastl's (2011) modification of Wilson's (1979) share auction model, which is commonly used in the literature (e.g. Cassola, Hortaçsu and Kastl, 2013; Hortaçsu, Kastl and Zhang, 2018; Allen, Kastl and Wittwer, 2022).

There are various ways to interpret the bidding costs, including the physical and time costs of submitting bids or of fine-tuning them.[30] They characterise one set of models that I consider, which is defined as follows.

Definition 1 (Model $S(\rho)$) *Model $S(\rho)$ is defined as the model which consists of Assumptions 3.1.1-3.1.4, with $F^{S(\rho)} = c_i(K_i)$, where $c_i(K_i + 1) \geq c_i(K_i) \geq 0 \, \forall K_i \in [1, \bar{K} - 1]$, and the risk aversion parameter is equal to ρ.*

The second of the two alternative assumptions on the friction F^m characterises a "truthful bidding" model (defined as follows), in which all bidders optimally submit bids which

[30] See Chapter 2 for discussion.

approximately correspond to their true marginal values. Specifically, each bidder faces an arbitrarily large cost of submitting a bid function which differs from the greatest integer function of their true marginal valuation function (i.e. for all quantities, the largest integer that is weakly less than their marginal value).

Definition 2 (Model T) *Model T is defined as the model which consists of Assumptions 3.1.1-3.1.4, with $F^T = \infty$ if $\sigma_i(\theta_i)$ differs from the greatest integer function of bidder i's marginal valuation function (i.e. for all quantities, the largest integer that is weakly less than bidder i's marginal value) for any $\theta_i \in \Theta_i$, and is zero otherwise, and $\rho \geq 0$.*[31]

The solution concept is a group-symmetric Bayesian Nash Equilibrium, in which bidders within the same group who receive the same signal adopt the same strategy. 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

$$\sigma_i(\theta_i) \in \underset{\sigma_i(\theta_i) \in \mathcal{F}_i}{\arg\max} V_i(\theta_i)$$

3.2 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 behaviour 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 she demands 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. 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 her marginal value so that $\mu_{i,k} = 0 \; \forall k \in \{1, ..., K_i\}, i \in \mathcal{N}$ in Model T. This is the case regardless of the bidder's degree of risk aversion.

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 her

[31] The degree of risk aversion does not affect the bidder's optimal strategy in Model T, which is the focus of analysis.

expected utility by unilaterally deviating from her prescribed strategy, σ_i, given her beliefs. Each bidder's strategy is multi-dimensional—for each auction, and for each type $\theta_i \in \Theta_i$, the bidder must choose the number of steps to submit, and for each step, a bid price and quantity to demand—so there are many possible ways in which to deviate.

The main identifying condition 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 her strategy constant.[32] Proposition 1 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 is the change in expected utility from winning a marginal unit at step k. The second effect is the change in expected utility due to the change in the expected payment that the bidder makes for the units that she wins if her demand affects the auction price. The tradeoff determines the optimal bid shading at the quantity the bidder demands at step k in equilibrium, i.e. $\mu_{i,k} = v^{g,h}(q_{i,k}) - b_{i,k}$. Equation 1 rules out a profitable deviation by setting these two effects to be equal.

The left-hand side of Equation 1 captures the first effect, which is the change in the bidder's expected utility holding the distribution of the auction price constant. It is helpful to split the effect by the possible realisations of the auction price. If the price is strictly below the bidder's bid at step k, but strictly above her bid at step $k+1$, the deviation increases her allocation by one unit, increasing her expected utility by the difference between her marginal value and the auction price (scaled by her total utility if she is risk averse). If, instead, the price equals her bid at step k, she is rationed, and her deviation only changes her allocation if another bidder is *also* rationed at step k, i.e. there is a "tie". (If the bidder's bid is the only one at price $b_{i,k}$, she is allocated the entire quantity supplied on the margin so her allocation is independent of the quantity that she demands.) Similarly, if the price equals her bid at step $k+1$, she is rationed and her deviation increases the quantity she is allocated only in the event of a tie, because the deviation increases the quantity she demands strictly above the auction price and correspondingly reduces the quantity she demands on the margin. If the auction price is either strictly above her bid at step k or strictly below her bid at step $k+1$, her deviation has no impact on the outcome of the auction, including on her own allocation.

[32] Chapter 2 presents a set of additional necessary conditions for equilibrium in Model $S(0)$, which rule out profitable unilateral deviations in the bid prices submitted by a bidder, holding the rest of her strategy constant.

The right-hand side of Equation 1 captures the second effect, which is the impact on the bidder's expected utility of the change in the distribution of the auction price due to her deviation.

I first state this necessary condition for equilibrium in Model $S(\rho)$ (Proposition 1). To build intuition, I then discuss special cases of the model, which make clear the implications of risk aversion for the amount of bid shading in equilibrium in Model $S(\rho)$.

Proposition 1 (Necessary condition on quantity deviations in Model $S(\rho)$) *In Model $S(\rho)$ with $\rho \geq 0$, 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*

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

where $P^g = P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)$ and $y_i = y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)$; $(P^g = p \wedge Tie^g)$ 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 her strategy, is unchanged.

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

To understand this condition and the implications of the degree of risk aversion for bid shading in equilibrium of Model $S(\rho)$, it is helpful to consider the special case in which there

are no ties on the margin for bidder i. In this case, Equation 1 implies bid shading of

$$\mu_{i,k} = \frac{\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[e^{\rho P^g q_{i,k}} P^g \middle| b_{i,k} > P^g > b_{i,k+1}\right]}{\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} - b_{i,k} + \begin{cases} \Omega q_{i,k} \frac{\partial}{\partial \epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho = 0 \\ \Omega \frac{1}{\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} \frac{1}{\rho} \frac{\partial}{\partial \epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[e^{\rho \tilde{P}^g(\epsilon) q_{i,k}}\mathbb{I}\left(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho > 0 \end{cases} \tag{2}$$

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

To interpret this bid shading, I consider bidders who take the distribution of the auction price as given and those with market power in turn.

3.2.1 Price-taking bidders

In Model $S(\rho)$, if a bidder is sufficiently small relative to the market, marginally increasing the quantity that she demands at step k has no impact on the auction price (the last term of the right-hand side of Equation 2 is zero).

In the case of no ties on the margin, the deviation only changes her expected utility if the market clears at a price strictly below her bid price at step k, $b_{i,k}$, but strictly above her bid price at step $k+1$, $b_{i,k+1}$. In this range of auction prices, the bidder wins precisely the cumulative quantity that she demands at step k. The change in her expected utility from deviating is the expected utility of winning the marginal unit, conditional on the auction price being in this range.

If the bidder is risk neutral, she demands a quantity at step k which equates her marginal value for the good with the conditional expected auction price, i.e. $v^{g,h}(q_{i,k}) = \mathbb{E}\left[P^g \middle| b_{i,k} > P^g > b_{i,k+1}\right]$. This trades off the possibility of winning the unit at a price at which the bidder would like to win, $P^g \leq v^{g,h}(q_{i,k})$, with winning at a price at which she would prefer to lose, $v^{g,h}(q_{i,k}) < P^g$. In equilibrium, her demand for the marginal unit is equivalent to accepting a lottery with an expected payoff of zero. This implies that she bids above her marginal value at the quantity she demands at step k so her bid shading is negative, $\mu_{i,k} < 0$.[33]

[33]Kastl (2011) discusses this somewhat surprising implication, which arises from the bidders' restricted strategy sets in Model $S(0)$. In contrast, a price-taking bidder would bid according to her marginal valuation

If the bidder is instead risk averse, she dislikes the possibility of a negative payoff and is only willing to accept the lottery if its expected payoff is sufficiently high. Risk aversion makes the bidder less aggressive in bidding for the marginal unit, which reduces the difference between the bid price she submits and her marginal value.

With sufficiently high risk aversion, any lottery with the possibility of a negative payoff is undesirable. The bidder therefore demands the maximum quantity at bid price $b_{i,k}$ for which the payoff from winning the marginal unit is non-negative. If the bidder may tie on the margin (so that $\mathbb{P}(P^g = b_{i,k} \wedge Tie^g) > 0$), this implies that the bidder will bid truthfully, because her payoff from winning the marginal unit is only non-negative if her bid is weakly below her marginal value at the quantity she demands. So, in the case of ties on the margin, the bidder will not shade her bid at all.

The case of no ties on the margin is more complex. When her bid is marginal, the quantity that the bidder demands does not determine her allocation so that a marginal increase in her demand has no impact on her utility. This means that the bidder's payoff from winning the marginal unit is non-negative when her bid is marginal, as well as when her bid is weakly below her marginal value. When the bidder is sufficiently risk averse, she demands the maximum quantity at bid price $b_{i,k}$ for which the payoff from winning the marginal unit is non-negative. With a continuous bid function, this implies that she will bid one basis point above her value (assuming the auction price distribution has full support). That is, the bidder's bid shading is -1.

3.2.2 Bidders with market power

More generally, a bidder has market power—the distribution of the auction price depends on the bids that she submits. Demanding a larger quantity weakly increases the price, which is paid for all units that she wins.

The last term of the right-hand side of Equation 2 captures the effect of this market power on the profitability of a marginal increase in the quantity demanded by bidder i at step k in Model $S(\rho)$. Holding the rest of her strategy constant, the deviation increases aggregate demand at prices between $b_{i,k}$ and $b_{i,k+1}$, which weakly increases auction prices in this range.

To isolate the effects of risk aversion, it is helpful to approximate the impact of the deviation

function in Wilson's (1979) classical model in which bidders submit continuous bid functions.

by a locally differentiable function.[34] Doing so, we can rewrite the last term of Equation 2 for $\rho \geq 0$ as:[35]

$$q_{i,k}\Omega\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{\partial P^g}{\partial q_{i,k}}\mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right] + C + D \tag{3}$$

where

$$C = q_{i,k}\Omega\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{\partial P^g}{\partial q_{i,k}}\mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right]\left(\frac{\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[e^{\rho P^g q_{i,k}}\middle| b_{i,k} > P^g \geq b_{i,k+1}\right]}{\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[e^{\rho P^g q_{i,k}}\middle| b_{i,k} > P^g > b_{i,k+1}\right]} - 1\right)$$

$$D = \frac{q_{i,k}\Omega Cov\left(e^{\rho P^g q_{i,k}}, \frac{\partial P^g}{\partial q_{i,k}}\middle| b_{i,k} > P^g \geq b_{i,k+1}\right)\mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)}{\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[e^{\rho P^g q_{i,k}}\middle| b_{i,k} > P^g > b_{i,k+1}\right]} \tag{4}$$

See Appendix A.2.

If the bidder is risk neutral, $C = D = 0$. In this case, 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 she demands at each price, in order to lower the auction price and reduce the total amount she must pay. The incentive to do so is increasing in the quantity that she demands.

If the bidder is risk averse, $\rho > 0$ so $C \leq 0$ and the sign of D depends on her beliefs over the distribution of the auction price.

The term C accounts for the fact that the marginal utility of a risk averse bidder is decreasing; she places greater weight on low-utility outcomes than she does on high-utility outcomes. In particular, she places disproportionate weight on the outcome where the original auction price equals her bid at step $k+1$ because her deviation would unambiguously reduce her utility: it would not change her allocation (assuming no ties) but might increase the auction price.[36] As risk aversion increases, $C \to -q_{i,k}\Omega\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{\partial P^g}{\partial q_{i,k}}\mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right]$, cancelling out the incentive of the risk neutral bidder to exploit market power (the first term of Equation 3), so that, if $D = 0$, the bidder does not shade her bid in order to exploit her market power.

[34] This is only an approximation because the bid functions of other bidders are also step functions.

[35] This uses the facts that $\mathbb{I}(b_{i,k} \geq \tilde{P}^g(\epsilon) \geq b_{i,k+1}) = \mathbb{I}\left(b_{i,k} \geq P^g \geq b_{i,k+1}\right)$ and that $\frac{\partial P^g}{\partial q_{i,k}} = 0$ for $P^g = b_{i,k}$ (because the bidder's deviation does not change the probability of $P^g \in [b_{i,k+1}, b_{i,k}]$ and the auction price is an increasing function of $q_{i,k}$) so the upper bound in the indicator function can be written strictly.

[36] In the knife-edge case that the bidder does not expect her bid at step $k+1$ to ever set the auction price, i.e. $\mathbb{P}(P^g = b_{i,k+1}) = 0$, then $C = 0$.

The term D is explained by the fact that a risk averse bidder prefers a less dispersed distribution of utility. If her market power is weaker at higher prices (so that $D < 0$), a reduction in the quantity that she demands reduces higher prices by less than it reduces lower prices. This deviation therefore increases the dispersion of the auction price distribution, and consequently increases the dispersion of the distribution of her utility. This weakens the incentive of the bidder to exert market power and reduces her bid shading at step k, relative to a comparable risk neutral bidder. Conversely, if her market power is stronger at higher prices, reducing the quantity she demands reduces the dispersion of the distribution of her utility, increasing her bid shading *ceteris paribus*.

The relationship between the auction price, P^g, and its derivative, $\frac{\partial P^g}{\partial q_{i,k}}$, depends on the bidder's beliefs over the distribution of other bidders' signals, as well as on the bid price that she submits.[37] And so, at moderate levels of risk aversion, D may be positive or negative and the relationship between bid shading and risk aversion is an empirical matter. Nonetheless, as risk aversion increases, $D \to 0$ for the same reasons why $C \to -q_{i,k}\Omega\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{\partial P^g}{\partial q_{i,k}}\mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right]$. Truthful bidding is therefore approximately optimal if the bidder is sufficiently risk averse.

4 Econometric method

4.1 Identification

In Model T, the marginal value of bidder $i \in \mathcal{N}$ at the quantity that she demands at step $k \in \{1, ..., K_i\}$ of her bid function is point identified by her 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 she demands at the steps of her bid functions from the observed bids, and hence to point identify her bid shading at these quantities.[38]

If bidder i does not expect to submit a bid at the same price as another bidder and therefore "tie" if her bid is marginal, her marginal value at the quantity she demands at step k can be readily estimated from Equation 2 (see Appendix B).

[37] The distribution of other bidders' signals implies a distribution of residual supply functions that the bidder faces. If the residual supply functions were vertical translations of one another and continuous, the relationship would be positive if the possible residual supply curves were convex (so $D > 0$), and negative if they were concave ($D < 0$).

[38] Equation 1 does not identify bidders' marginal values at quantities between the quantities demanded at the steps. In Chapter 2, I provide additional necessary conditions for equilibrium in Model $S(0)$ which can be used to set-identify bidders' full marginal valuation functions.

However, bidders frequently submit bids at the same prices as other bidders in the ILTR auctions because they can only submit bids in increments of whole basis points. This is especially the case for good A, for which the majority of bids (85.1% in June 2010 – May 2012) are submitted between 0 and 5bps. For goods A and B, respectively, 69% and 20% of bids in June 2010 – May 2012 are submitted at prices at which at least one other bidder bids in the same auction.

Allowing for ties makes identification more complex. 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 she is allocated by the product of the marginal quantity she demands at step k and the rationing coefficient.[39] This deviation therefore affects the bidder's utility differently in the case that her bid is marginal relative to the case that her bid is strictly above the auction price.

To estimate a bidder's marginal value at step k using Equation 1, we therefore require three further pieces of information beyond what is required in the case of no ties:

1. The rationing coefficient at step k, which determines the quantity the bidder wins when rationed at step k.
2. The bidder's marginal value at the quantity she wins when rationed at step k, which is the marginal benefit of winning a additional unit at step k when the bidder is risk neutral, *ceteris paribus*.
3. The bidder's marginal value function between the quantity she wins when rationed at step k and the quantity she demands at step k. When the bidder is risk averse, this impacts the bidder's overall net utility from the auction and therefore the marginal benefit of winning an additional unit, *ceteris paribus*.

Bidders beliefs about the size of the rationing coefficient can be estimated from the data (see following section), and I make two additional assumptions to recover the two remaining pieces of information.

The first assumption, Assumption 4.1, is that bidder i's marginal valuation function is flat around the quantity that she demands, so that her marginal value at the quantity that she

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

wins when rationed is the same as her marginal value at the quantity she demands. 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 $S(0)$.[40] Some natural alternative assumptions would achieve the same result as Assumption 4.1, e.g. that the marginal value function is piecewise linear. For bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$,

$$v^{g,h}(y_i, \theta_i) = v^{g,h}(q_{i,k}, \theta_i) \ \forall y_i \in \Big[q_{i,k-1} + r_{i,k}(q_{i,k} - q_{i,k-1}), q_{i,k} + r_{i,k+1}(q_{i,k+1} - q_{i,k})\Big).$$

The second assumption, Assumption 4.1, is that the rationing coefficient that bidder i expects to face at price $b_{i,k}$, conditional on both the strategy profile in Model $S(\rho)$ and there being a tie at the price $P^g = b_{i,k}$, is deterministic. This approximates a situation in which the bidder does not take the variation in her net utility from the auction into account when evaluating the marginal benefit of winning an additional unit at a auction price equal to her bid.[41]

The rationing coefficient for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}$ at step $k \in \{1, ..., K_i\}$ if she ties, denoted $r_{i,k} = \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{\partial y_i}{\partial q_{i,k}} \middle| P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i) = b_{i,k} \wedge Tie^g\right]$, is deterministic.

Under these assumptions, the necessary condition for bidders' strategies to be a Bayesian

[40] 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 she is allocated when rationed at step $k + 1$.

[41] We could relax Assumption 4.1 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.

Nash Equilibrium in Model $S(\rho)$ simplifies to

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}}\left(v^{g,h}(q_{i,k},\theta_i)-P^g\right)\Big|b_{i,k}>P^g>b_{i,k+1}\right]\mathbb{P}(b_{i,k}>P^g>b_{i,k+1})\\
&+\Bigg(e^{\rho(1-r_{i,k})(q_{i,k}-q_{i,k-1})v^{g,h}(q_{i,k},\theta_i)}e^{\rho b_{i,k}\left(q_{i,k-1}+r_{i,k}\left(q_{i,k}-q_{i,k-1}\right)\right)}\\
&\qquad\qquad\left(v^{g,h}\left(q_{i,k},\theta_i\right)-b_{i,k}\right)r_{i,k}\Bigg)\mathbb{P}(P^g=b_{i,k}\wedge Tie^g)\\
&+\Bigg(e^{-\rho r_{i,k+1}(q_{i,k+1}-q_{i,k})v^{g,h}(q_{i,k},\theta_i)}e^{\rho b_{i,k+1}\left(q_{i,k}+r_{i,k+1}\left(q_{i,k+1}-q_{i,k}\right)\right)}\\
&\qquad\qquad\left(v^{g,h}\left(q_{i,k+1},\theta_i\right)-b_{i,k+1}\right)\left(1-r^g_{i,k+1}\right)\Bigg)\mathbb{P}(P^g=b_{i,k+1}\wedge Tie^g)\\
=&\begin{cases}
q_{i,k}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(b_{i,k}\geq\tilde{P}^g(\epsilon)\geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho=0\\
\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho\tilde{P}^g(\epsilon)q_{i,k}}e\mathbb{I}\left(b_{i,k}\geq\tilde{P}^g(\epsilon)\geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho>0
\end{cases}
\end{aligned}
\tag{5}
$$

See Appendix A.3.

Equation 5 can be solved numerically to recover bidders' marginal values, by calibrating the risk aversion parameter and using the bidding data to estimate bidders' beliefs over the distribution of the auction price and the impact of her own bids on this distribution.

It is not possible to identify the degree of risk aversion from Equation 5, so I test the relative performance of models which differ in this parameter in Section 4.3. I also estimate bidders' marginal values assuming no ties using Equation 2 (see Appendix B).

I turn next to the estimation method for the remaining components of Equation 5.

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 ease of notation, I suppress time subscripts, except for when describing the resampling method for estimating a bidder's beliefs and the sample of other bidders used.

For Model T, the marginal value of bidder $i \in \mathcal{N}$ at the quantity she demands at step k of her bid function is trivially estimated by her 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, Perrigne and Vuong (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. This strategy profile is determined by bidders' common prior of the number of potential bidders and the signal distribution, and the mapping from signals to strategies. Guerre, Perrigne and Vuong (2000) show that, assuming equilibrium, the realised 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.

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.[42] 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 her 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.[43]

The resulting empirical distribution of auction prices for good g can be used to consistently estimate the expectation terms in Equation 5. The derivative term is estimated using a numerical derivative, as described in Kastl (2011).[44] Given these estimates, the marginal

[42] A bidder who submits bids on both goods is treated as two separate bidders.

[43] The sample of other bidders is finite, so only a finite number of signal realisations is ever (implicitly) observed and the empirical distribution of auction prices may not have full support. This creates issues for calculating the conditional expectation terms and the derivative of the auction price. The former may not be identified and the latter may be numerically unstable. To address this, I use kernel density estimation to smooth the auction price distribution and ensure full support, discretised at integer prices.

[44] The quantity demanded by bidder i at step k is perturbed by $\epsilon > 0$, where ϵ is small, holding the rest of

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 $q_{i,k+1} = 0$). I solve for $v^{g,h}(q_{i,k}, \theta_i)$ numerically.

Sample of other bidders 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 neighbours of the same term, and stratify by the grouping defined in Section 2.3.[45,46] This aims to maintain consistency in the economic environment and therefore signal distribution, while increasing precision by pooling bids from neighbouring auctions.

4.3 Testing procedure

Models $S(\rho)$, $\rho \geq 0$, 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 behaviour.

Ideally we would compare the bid shading implied by a candidate model to bid shading in the correct model of behaviour – if the discrepancy is only small, the candidate model is a good approximation. Because the actual amount of bid shading is unobserved, I follow Backus, Conlon and Sinkinson's (2021) approach, which uses the exclusion restrictions proposed by Berry and Haile (2014) to mimic this comparison. For a set of instruments that do not

her strategy fixed. The auction price distribution is then recalculated for the bidder. The derivative term is the difference in expectations under the perturbed and unperturbed distributions, scaled by the size of the perturbation.

[45]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 neighbours, rather than three or more.

[46]Under the assumption of i.i.d. signals within group, as the sample becomes large, there is no reason to exclude observation (i, t) from the sample of other bidders, $\mathcal{O}_{i,t}$. Observation (i, t) represents the best response of one possible signal realisation of another bidder. However, with a finite sample including observation (i, t) in $\mathcal{O}_{i,t}$ artificially increases the estimated probability that bidder i ties on the margin in auction t, biasing the estimator of the bidder's beliefs. To address this, $(i, t) \notin \mathcal{O}_{i,t}$.

affect bidders' true values, the bid shading implied by a candidate model is projected on the instruments and implicitly compared to the correct model's bid shading projected on the instruments. Any misspecification of the candidate model will cause a discrepancy between these two predictions. The bid shading in my context is analogous to the markups implied by alternative models of firm conduct in product markets that Backus, Conlon and Sinkinson (2021) analyse.

The test is a pairwise comparison of the fit of alternative models of behaviour.[47] The test assesses only the relative performance of two models, and so is robust to their misspecification.[48] This is particularly relevant in the empirical auction literature, in which a model is often required to recover unknown parameters in order for further analysis, and is necessarily a simplification of reality.[49] The question we ask is, among a set of candidate models of behaviour, which is the best approximation in our setting.

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 vice versa.

The test is based on the relationship between the model-implied marginal values 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 her value. I define the true marginal value of bidder i at step k of her bid function in auction t by $v_{i,k,t}$, a function of a set of exogenous

[47] It is applicable to non-nested alternatives, which is relevant in our setting to compare Model $S(0)$ with Model $S(\rho)$, $\rho > 0$.

[48] In Chapter 2, I analyse an external measure of fit of Model $S(0)$, by estimating lower bounds on the bidding costs implied by the model.

[49] For example, in the risk-neutral case, 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 $S(\rho)$ is a better approximation. For example, if bidders believe they face much stronger competition than in reality, Model T might be the better approximation.

characteristics, $\boldsymbol{x}$, and an additively separable residual value, $\omega_{i,k,t}$. The critical assumption is that condition 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 her 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 her 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. For simplicity let $\boldsymbol{x}$ be a constant so that we can ignore the distinction between $v_{i,k,t}$ and $\omega_{i,k,t}$ in this example. Suppose that Model T is correct and Model $S(0)$ is incorrect (but the reasoning holds for each pair of models). 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 her true value is equal to her 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 correlation 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 $S(0)$: $Cor(v^T_{i,k,t},z_{i,k,t})=Cor(v_{i,k,t},z_{i,k,t})=0$ and $Cor(v^{S(0)}_{i,k,t},z_{i,k,t})=Cor(b_{i,k,t},z_{i,k,t})+Cor(\mu^{S(0)}_{i,k,t},z_{i,k,t})=Cor(v_{i,k,t},z_{i,k,t})+Cor(\mu^{S(0)}_{i,k,t},z_{i,k,t})=Cor(\mu^{S(0)}_{i,k,t},z_{i,k,t})\neq 0$.[50] The crux of the test is that the correct model would have zero correlation and a larger correlation implies a worse model fit. In general, neither model will be correct and, broadly speaking, the test compares the relative magnitude of the correlations.

The steps of the test modified from Backus, Conlon and Sinkinson (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, $v^m_{i,k,t}=b_{i,k,t}+\mu^m_{i,k,t}$.

[50] It is straightforward to see the converse. Suppose that Model $S(0)$ is correct and Model T is incorrect. The bidder's true marginal value is therefore equal to the marginal value implied by Model $S(0)$, $v_{i,k,t}=v^{S(0)}_{i,k,t}=b_{i,k,t}+\mu^{S(0)}_{i,k,t}$. So the marginal value implied by Model T is $v^T_{i,k,t}=b_{i,k,t}=v^{S(0)}_{i,k,t}-\mu^{S(0)}_{i,k,t}=v_{i,k,t}-\mu^{S(0)}_{i,k,t}$, and the correlation between the model-implied marginal values and a valid instrument, $z_{i,k,t}$, will be zero for Model $S(0)$ and non-zero for Model T: $Cor(v^{S(0)}_{i,k,t},z_{i,k,t})=Cor(v_{i,k,t},z_{i,k,t})=0$ and $Cor(v^T_{i,k,t},z_{i,k,t})=-Cor(\mu^{S(0)}_{i,k,t},z_{i,k,t})\neq 0$.

(b) Estimate $v^m_{i,k,t}$ as a function, $h(\boldsymbol{x})$, of the set of exogenous characteristics $\boldsymbol{x}$, which are specified below, 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 the instruments, $\boldsymbol{z}$, and the function, $g(.)$, are specified below. Obtain the predictions, $\Delta\hat{\mu}_{i,k,t}$.

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.

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,[51] 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.[52] 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 observation), conditional on the bidder's size and on the auction.[53,54]

Bid shading is affected by bidders' beliefs about the strength of competition that they face at each price. There are many potential ways to measure this 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 herself (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 she did not bid at all, and moments of this distribution; and moments of the distribution of bids excluding the bidder's own bids. Specifically, the set of instruments, $\boldsymbol{z}$, for bidder i in auction t, which are assumed to satisfy the exogeneity condition, $\mathbb{E}\left[\omega_{i,k,t}|\boldsymbol{x}, \boldsymbol{z}\right] = 0$, includes:

[51] These are equivalent to time dummies because up to one auction is held per day.

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

[53] The quantities demanded at the bids are excluded from $\boldsymbol{x}$ because, if Model $S(\rho)$ is the correct model of behaviour, 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 Chapter 2).

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

- For goods $g = \{A, B\}$, number of bidders, number of small bidders, and number of large bidders, each excluding bidder i, in auction t[55]
- The average bid in auction t, excluding the bids of bidder i
- For $P_{i,t}^{g}$, defined by the clearing price of good g in auction t, which is calculated using the resampling procedure described in Section 4.2 but excluding bidder i's bids when calculating the equilibrium distribution of the auction price,
  - The conditional expected price for good A, $\mathbb{E}\left[P^A \leq X\right]$ for $X = [1, 5, 10]$
  - The conditional expected price for good B, $\mathbb{E}\left[P^B \leq X\right]$ for $X = [6, ..., 40]$
  - The conditional expected price differences for B: $(X - \mathbb{E}\left[X - 5 \leq P^B \leq X\right])$ for $X = [6, ..., 40]$ and $(X - \mathbb{E}\left[X - 10 \leq P^B \leq X\right])$ for $X = [6, ..., 40]$[56]
- All of the above set interacted with good-term-specific dummies

Under the assumption made throughout the analysis that bidders have private values (see Section 3), these measures of competition do not impact the bidder's own marginal valuation function, and so can be included in the set of instruments, $\boldsymbol{z}$.

This implies a large number of potential instruments (equal to 460). Following Carrasco (2012), Conlon (2017) and Backus, Conlon and Sinkinson (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 set to include in $\boldsymbol{z}$. The relationship between the instruments—measures of competition—and bid shading, are non-linear (for Model $S(\rho)$, this is shown by Equation 1). To capture these non-linearities while preventing overfitting, the functions $h(.)$ and $g(.)$ are estimated using a random forest.

I restrict the set of observations to bids strictly above the reserve prices.[57] 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.

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

[56] The bounds of the conditional expectations represent sensible points on the distribution of bids (see Table 1).

[57] 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.

5 Results

5.1 Model fit

I run the test for pairs of models among Models $S(\rho)$ for $\rho \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 1\}$ and T. Results are reported in Table 3. Each entry shows the test statistic for the row being Model 1 and the column being Model 2, so that a negative entry indicates that the row model has better model fit than the column model (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%.

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 3 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 between goods.

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 must be whole basis points), 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 conventional model.

In Chapter 2, I analyse an external measure of fit of Model $S(0)$. The measure is the size of the bidding costs required to explain the data under the assumptions of the model. Larger

bidding costs suggest a worse model fit because the physical costs of submitting bids are trivial in the BoE's ILTR auctions.[58] I estimate a lower bound on the bidding costs by the marginal benefit of submitting an additional bid.[59] For good A, the estimated bidding costs are trivial, suggesting that the model characterises bidding behaviour well by this measure. This is not the case for good B, at least for the 67% of bidders for good B who submit only one bid. This complements Chapter 1, which finds an improvement on Model $S(0)$.

Table 3: Testing results

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	-1.982	-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*							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	0.215	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*							
Model 1	Model 2						
	$S(0)$	$S(0.1)$	$S(0.2)$	$S(0.3)$	$S(0.4)$	$S(0.5)$	$S(1)$
T	-2.574	-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

Each entry shows the test statistic for the row being Model 1 and the column being Model 2, so that a negative entry indicates that the row model has better model fit than the column model (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%.

[58] Interpreting the costs as cognitive costs does not seem to improve the model fit (see Appendix 2 of Chapter 2).

[59] For a bidder who submits K bids, the bidding cost of the $(K+1)$th bid must be larger than the marginal benefit of "fine-tuning" her bid function, which is equal to the difference in her expected utility from submitting what would be the optimal $(K+1)$-step bid function and what would be the optimal K-step one. This marginal benefit of fine-tuning, which provides a lower bound on the bidding cost of the $(K+1)$th bid.

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

I consider two ways to rationalise 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' behaviour, 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.[60,61] 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 be observationally equivalent. This explains the results of the pairwise test in Table 3. 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. These results do not imply that bid shading is necessarily the same in the two models, but rather that the models explain the data equally well.

To examine the degree of risk aversion required for bidders to bid truthfully in Model $S(\rho)$, Figure 3 plots the absolute value of bid shading implied by Model $S(\rho)$ against the log of

[60] In a more general model, a bidder might choose to avoid the cost of precisely calculating the optimal strategy by choosing an approximation to it. In Wilson's (1979) classical model, truthful bidding is optimal for small bidders anyway but larger bidders optimally submit a bid function below their marginal valuation function in order to exploit market power, in which case, a small amount of bid shading might be a better approximation than truthful bidding. This is not the case in Model $S(\rho)$, in which the optimal amount of bid shading can be positive or negative.

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

the risk aversion parameter, $log(\rho)$.[62] It shows that almost all bidders 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, aggregate bidder surplus per auction, assuming bidders bid truthfully, is £220,461 in the 3-month auctions and £326,687 in the 6-month auctions on average (weighted by the quantities allocated per auction). Under the CARA specification, a bidder with $log(\rho) = 4$ facing a 50-50 bet to lose the entire surplus of £220,461 or gain £X would accept if $X >$ 220,993 in the 3-month auctions.[63] In the 6-month auctions, the corresponding loss/gain is £326,687/£329,034. For 50-50 bets where the size of the loss corresponds to individual bidder surplus (which is of course lower than the aggregate), the minimum gain to take the bet would be relatively smaller.[64] This suggests that a fairly low degree of risk aversion is needed for truthful bidding to be optimal in the ILTR auctions.

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 estimated for bidders 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 $log(\rho) = 5.04$ and $log(\rho) = 9.83$, respectively.[65] In US Treasury auctions, Boyarchenko, Lucca and Veldkamp (2021) estimate a parameter which is on average approximately equivalent to $log(\rho) = 4.48$.[66] Using any

[62]The results show that a small amount of risk aversion has only a very small effect on optimal bid shading—bid shading is roughly constant as a function of ρ in the range $log(\rho) < -2.5$ ($\rho < 0.082$). This is largely driven by the fact that the optimal amount of bid shading for all bidders depends on the value of a marginal unit (see Section 3.2.1), and the rate at which this value declines as a function of ρ is increasing if ρ is small.

[63]In the model, the maximum quantity supplied in an auction is normalised to one, so that the risk aversion parameter, ρ, measures the rate at which a bidder's marginal utility decreases when her net utility from the auction (or her "wealth") 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 £ would be $(0.2 \times 10^{-9} \times \rho)$ and $(0.4 \times 10^{-9} \times \rho)$, respectively.

[64]For example, a bidder with $log(\rho) = 4$ facing a 50-50 bet to lose £50,000 and gain £X would accept if $X >$ 50,027 in the 3-month auctions and if $X >$ 50,055 in the 6-month auctions.

[65]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 (the 1998 auctions were held in francs (FRF) and converted to euros at a rate of 6.55957 francs per euro). I convert these at the average EUR/GBP exchange rate in May 1998 – December 2000 of 1.55 (where the EUR/GBP exchange rate in 1998 is calculated by (1/6.55957)FRF/GBP). The average maximum supply in my study is £4.17 billion, so Armantier and Sbaï's (2006) estimates imply $\rho = \frac{(6.907 \times 10^{-6})}{1.55} \times 4.17 \times 10^9$ and $\rho = \frac{(5.732 \times 10^{-8})}{1.55} \times 4.17 \times 10^9$, respectively.

[66]Boyarchenko, Lucca and Veldkamp (2021) estimate a CARA parameter of 366.61 given a total auction supply normalised to one, which implies $\rho = 366.61\hat{Q}$, where $\hat{Q}$ is the ratio of the average maximum supply

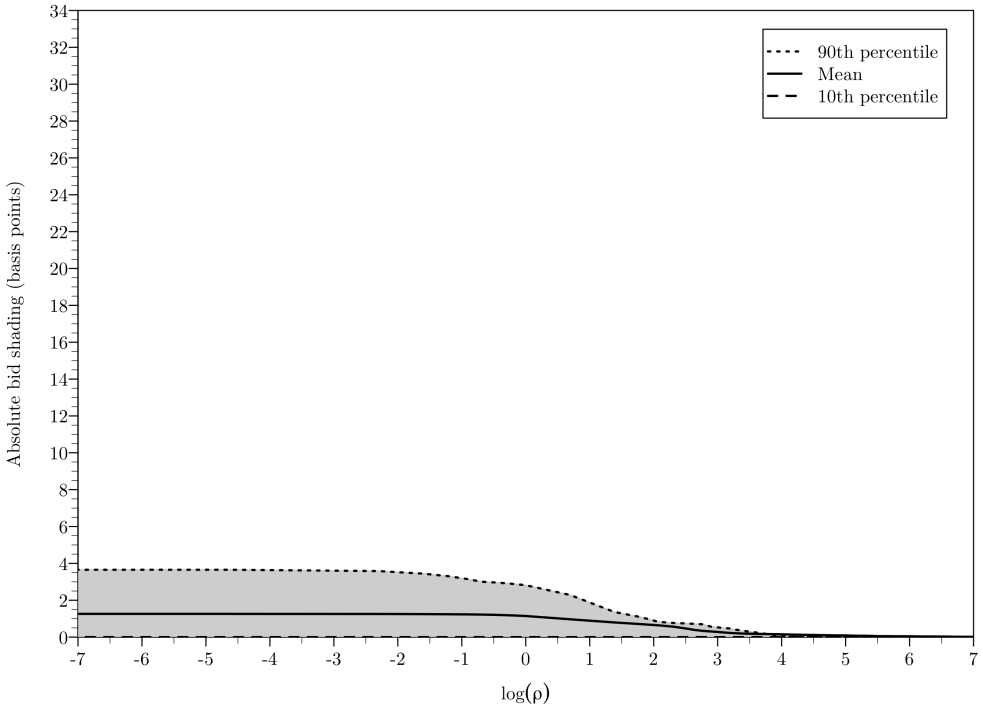
of these values in Model $S(\rho)$ in my setting would suggest that truthful bidding can be rationalised by the fact that bidders are risk averse.

Nonetheless, the degree of risk aversion estimated in auctions varies across studies (see Vasserman and Watt (2021) for a survey), and the appropriate degree of risk aversion in our setting depends on the interpretation. One interpretation is that the incentives of the manager tasked with bidding on behalf of the participant lead to risk-averse behaviour within a principal-agent framework. For example, 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. An alternative interpretation is that the financial institutions themselves are risk averse, given the BoE's auctions were introduced to respond to "the new demands for liquidity insurance that [the financial crisis] engendered" (Fisher, 2011a, p.15), suggesting bidders would be willing to pay higher prices to reduce the uncertainty in their allocations. However, the sample period turned out to be relatively stable and participants' bids did not appear to reflect acute liquidity needs so a principal-agent interpretation seems more appropriate.

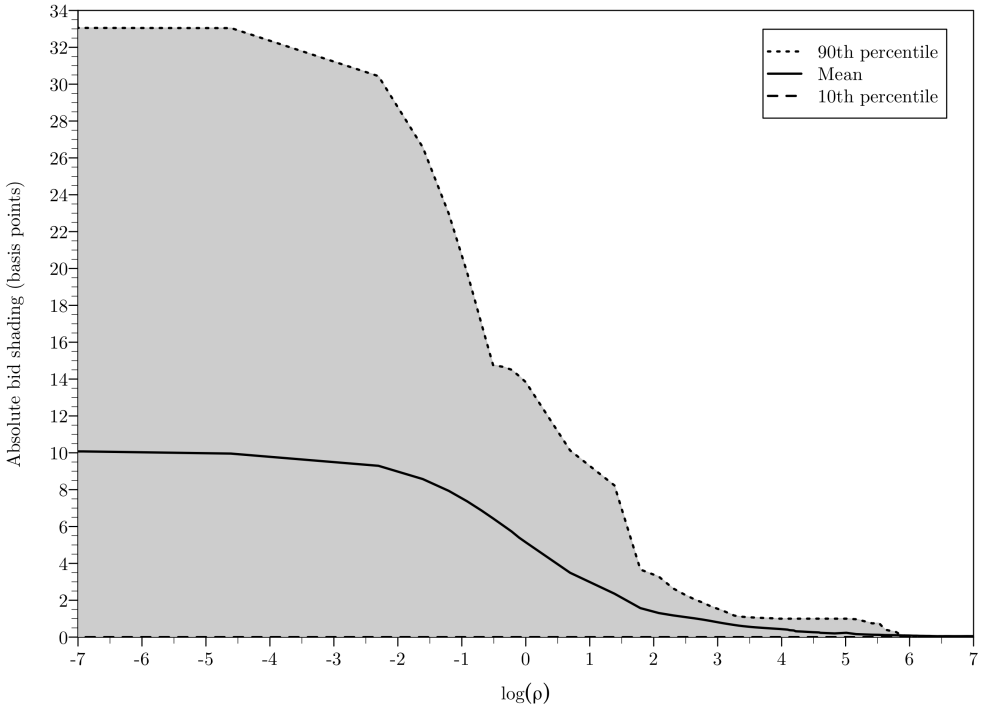
in my model to the average issuance in their study. The average maximum supply in my study is £4.17 billion; their sample includes US Treasury auctions of 2-, 3-, 5-, 7-, and 10-year notes in September 2004 – June 2014, of which the average issuance was £17.3 billion (converted from US dollars at the exchange rate on issuance date), which implies $\hat{Q} = 0.24$.

Figure 3: Absolute value of bid shading in Model $S(\rho)$ for $log(\rho) \in [-10, 10]$, June 2010 – May 2012 (observations are weighted by the quantities demanded)

(a) Good A



(b) Good B



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-maximising behaviour.

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 bid 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 bid truthfully and are risk neutral.[67,68] 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.[69]

[67] 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 is analogous to the approach described in Chapter 2.

[68] Allowing for risk aversion would reduce the incremental profit from what otherwise would be the optimal strategy. This would reduce the estimated lower bound on the cost of sophistication that explains why bidders bid truthfully.

[69] One reason is that Chapter 2 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. While these benefits are estimated under slightly different assumptions, namely that the bidder is bidding optimally rather than truthfully, this suggests that the results would not change much if we estimated a

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 rationalise 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).

6 Conclusion

The 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 behaviour which is more conventional in the literature—bidders choose to submit bids which may differ from their true values in order to maximise their own expected surplus—the measured efficiency of, and information gleaned from, the auctions would change.[70]

I therefore compare the relative performance of these alternative models of behaviour. I find that bidding behaviour 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, Conlon and Sinkinson (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 in similar contexts.

three-step bid function. 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.

[70] Cassola, Hortaçsu and Kastl (2013) disaggregate the change in bidding in the 2007 subprime market crisis in the European Central Bank's repo auctions into a change in valuations and a strategic response.

Alternatively, truthful bidding can be rationalised by a cost of determining what otherwise would be 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 behaviour—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.

Appendices

A Proof of Proposition 1

Kastl (2011) proves Proposition 1 for $\rho = 0$. I follow his approach for the case of $\rho > 0$.

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 her 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 her strategy, is unchanged.

For $\rho > 0$, we aim to evaluate the limit:

$$\lim_{\epsilon \to 0} \frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1 - e^{-\rho U_i(y_i(\epsilon), P^g(\epsilon))}\right)\right] - \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1 - e^{-\rho U_i(y_i(0), P^g(0))}\right)\right]}{\epsilon} \tag{6}$$

where I write $U_i(.|\theta_i)$ as $U_i(.)$ for convenience.

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, which partitions the state space as follows

$$\begin{aligned}
\theta_{1k}(x) &= \{\boldsymbol{\Theta_{-i}} : b_{i,k+1} < P^g(x) \leq b_{i,k},\ y_i(x) = q_{i,k} - x\} \\
\theta_{2k}(x) &= \left\{\boldsymbol{\Theta_{-i}} : P^g(x) = b_{i,k},\ y_i(x) = y_i^{RAT}(q_{i,k} - x - q_{i,k-1}) : q_{i,k-1} < y_i^{RAT} < q_{i,k} - x\right\} \\
\theta_{3k}(x) &= \left\{\boldsymbol{\Theta_{-i}} : P^g(x) = b_{i,k+1},\ y_i(x) = y_i^{RAT}(q_{i,k+1} - q_{i,k} + x) : q_{i,k} - x < y_i^{RAT} < q_{i,k+1}\right\} \\
\theta_{4k}(x) &= \{\boldsymbol{\Theta_{-i}} : b_{i,k} < P^g(x),\ y_i(x) \leq q_{i,k-1}\} \\
\theta_{5k}(x) &= \{\boldsymbol{\Theta_{-i}} : P^g(x) < b_{i,k+1},\ q_{i,k+1} \leq y_i(x)\}
\end{aligned}$$

for $x \in \{0, \epsilon\}$, where $\epsilon > 0$ and $y_i^{RAT}(q)$ is the quantity that bidder i is allocated when rationed, given the quantity that i demands at the margin is q.

The set $\theta_{1k}(x)$ corresponds to the outcomes in which bidder i is allocated precisely the quantity she demands at step k; $\theta_{2k}(x)$ and $\theta_{3k}(x)$ correspond to the outcomes in which bidder i is rationed at step k and step $k+1$, respectively (the distinction between tying with another bidder or not when rationed is made later in the proof); $\theta_{4k}(x)$ and $\theta_{5k}(x)$ correspond

to the outcomes in which bidder i is allocated strictly less and strictly more, respectively, than the quantity she demands at step k. We then define the following subsets of θ_{2k} and θ_{3k} as

$$\omega_{1k}(\epsilon) = \{\boldsymbol{\Theta}_{-\boldsymbol{i}} : \boldsymbol{\Theta}_{-\boldsymbol{i}} \in \theta_{2k}(0) \cap \theta_{1k}(\epsilon)\}$$
$$\omega_{2k}(\epsilon) = \{\boldsymbol{\Theta}_{-\boldsymbol{i}} : \boldsymbol{\Theta}_{-\boldsymbol{i}} \in \theta_{2k}(0) \cap \theta_{3k}(\epsilon)\}$$
$$\omega_{3k}(\epsilon) = \{\boldsymbol{\Theta}_{-\boldsymbol{i}} : \boldsymbol{\Theta}_{-\boldsymbol{i}} \in \theta_{1k}(0) \cap \theta_{3k}(\epsilon)\}$$

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 her 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:

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i(\epsilon),P^g(\epsilon))}\right)\right]-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left(\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i(0),P^g(0))}\right)\right]\\
=&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\theta_{1k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\theta_{1k}(0)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(\epsilon)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\theta_{3k}(\epsilon)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\right);\theta_{3k}(0)\right]\\
=&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\theta_{1k}(0)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\theta_{3k}(0)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\right);\theta_{3k}(0)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\theta_{1k}(0)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\omega_{1k}(\epsilon)\right]\\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\omega_{2k}(\epsilon)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\omega_{3k}(\epsilon)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\omega_{2k}(\epsilon)\right]\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\omega_{1k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\omega_{3k}(\epsilon)\right] \quad (7)
\end{aligned}
$$

where the last line follows from the facts that

$$\begin{aligned}
\mathbb{P}\left(\theta_{1k}(\epsilon)\right) &= \mathbb{P}\left(\theta_{1k}(0)\right) + \mathbb{P}\left(\omega_{1k}(\epsilon)\right) - \mathbb{P}\left(\omega_{3k}(\epsilon)\right) \\
\mathbb{P}\left(\theta_{2k}(\epsilon)\right) &= \mathbb{P}\left(\theta_{2k}(0)\right) - \mathbb{P}\left(\omega_{1k}(\epsilon)\right) - \mathbb{P}\left(\omega_{2k}(\epsilon)\right) \\
\mathbb{P}\left(\theta_{3k}(\epsilon)\right) &= \mathbb{P}\left(\theta_{3k}(0)\right) + \mathbb{P}\left(\omega_{3k}(\epsilon)\right) + \mathbb{P}\left(\omega_{2k}(\epsilon)\right)
\end{aligned}$$

Note that

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{2k}(0)\right] - \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{2k}(\epsilon)\right] \\
&-\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{1k}(\epsilon)\right] - \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{2k}(\epsilon)\right] = 0
\end{aligned} \tag{8}$$

and

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{3k}(0)\right] - \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{3k}(\epsilon)\right] \\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{3k}(\epsilon)\right] + \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{2k}(\epsilon)\right] = 0
\end{aligned} \tag{9}$$

Summing Equations 7-9, adding and subtracting

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{1k}(0)\right], \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\theta_{2k}(0)\right], \\
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\theta_{3k}(0)\right], \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k})}\right);\omega_{1k}(\epsilon)\right], \\
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k})}\right);\omega_{2k}(\epsilon)\right], \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k+1})}\right);\omega_{3k}(\epsilon)\right], \\
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k+1})}\right);\omega_{2k}(\epsilon)\right],
\end{aligned}$$

and collecting terms, we therefore aim to evaluate

$$
\begin{aligned}
&\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i(\epsilon),P^g(\epsilon))}\right)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left(\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i(0),P^g(0))}\right)\right]\right)\right)\\
=&\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\theta_{1k}(0)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{1k}(0)\right]\right)\right)\\
&+\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\right.\right.\\
&\left.\left.-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\right)\right)\\
&+\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\theta_{3k}(0)\right]\right.\right.\\
&\left.\left.-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\right);\theta_{3k}(0)\right]\right)\right)\\
&+\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right)-\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\bigcup_{j=1}^{3}\theta_{jk}(0)\right]\right)\right)\\
&+\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k})}\right);\omega_{1k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\omega_{1k}(\epsilon)\right]\right.\right.\\
&\left.\left.+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k})}\right);\omega_{2k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\omega_{2k}(\epsilon)\right]\right.\right.\\
&\left.\left.+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\omega_{3k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k+1})}\right);\omega_{3k}(\epsilon)\right]\right.\right.\\
&\left.\left.+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\omega_{2k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k+1})}\right);\omega_{2k}(\epsilon)\right]\right.\right.\\
&\left.\left.+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\omega_{1k}(\epsilon)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{1k}(\epsilon)\right]\right.\right.\\
&\left.\left.-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\omega_{3k}(\epsilon)\right]+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\omega_{3k}(\epsilon)\right]\right)\right)
\end{aligned}
\tag{10}
$$

Recalling the definition of $U_i(y_i(x), P^g(x))$, we have 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

$$\begin{aligned}
&\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k}-\epsilon,P^g(\epsilon))}\right);\theta_{1k}(0)\right]-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\theta_{1k}(0)\right]\right)\\
&\qquad=\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho U_i(q_{i,k},P^g(0))}\left(v^{g,h}(q_{i,k},\theta_i)-P^g(0)\right);\theta_{1k}(0)\right]\\
&\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\right.\\
&\qquad\left.-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k}-q_{i,k-1}),b_{i,k})}\right);\theta_{2k}(0)\right]\right)\\
&\qquad=\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho U_i(y_i^{RAT}(q_{i,k}-q_{i,k-1}),b_{i,k})}\left(v^{g,h}(y_i^{RAT}(q_{i,k}-q_{i,k-1}),\theta_i)-b_{i,k}\right);\theta_{2k}(0)\right]\\
&\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}+\epsilon),b_{i,k+1})}\right);\theta_{3k}(0)\right]\right.\\
&\qquad\left.-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\right);\theta_{3k}(0)\right]\right)\\
&\qquad=\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho U_i(y_i^{RAT}(q_{i,k+1}-q_{i,k}),b_{i,k+1})}\left(v^{g,h}(y_i^{RAT}(q_{i,k+1}-q_{i,k}),\theta_i)-b_{i,k+1}\right);\theta_{3k}(0)\right]
\end{aligned}$$

Now consider the fourth limit. We can show that the expectation term is continuous in ϵ by first partitioning $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\bigcup_{j=1}^{3}\theta_{jk}(0)\right]$, into

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right)\middle|\theta_{1k}(0)\right]\mathbb{P}\left(\theta_{1k}(0)\right)\\
&+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k})}\right)\middle|\theta_{2k}(0)\right]\mathbb{P}\left(\theta_{2k}(0)\right)+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},b_{i,k+1}}\right)\middle|\theta_{3k}(0)\right]\mathbb{P}\left(\theta_{3k}(0)\right).
\end{aligned}$$

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 $\rho>0$, it follows that the exponential is bounded, i.e. $\left[e^{\rho U_i(q_{i,k},P^g(\epsilon))}\middle|\theta_{1k}(0)\right]\in[0,1]$. Therefore by the same reasoning as Kastl's (2011) Lemmas A3 and A4, $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right)\middle|\theta_{1k}(0)\right]$ is continuous in ϵ at $\epsilon=0$ for a.e. $\theta_i\in\Theta_i$, is of bounded variation and satisifes 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, we have

$$\lim_{\epsilon\to 0}\left(\frac{1}{\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right)-\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(0))}\right);\bigcup_{j=1}^{3}\theta_{jk}(0)\right]\right)\right)$$
$$=\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(q_{i,k},P^g(\epsilon))}\right);\bigcup_{j=1}^{3}\theta_{jk}(0)\right]\right)\Bigg|_{\epsilon=0}$$

And finally, we can express the expectations in the last limit in terms of their conditional expectations, i.e.

$$\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y^g_{i,k}(x),P^g(x))}\right);\omega_{jk}(\epsilon)\right]=\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\left[\frac{1}{\rho}\left(1-e^{-\rho U_i(y^g_{i,k}(x),P^g(x))}\right)\middle|\omega_{jk}(\epsilon)\right]\mathbb{P}\left(\omega_{jk}(\epsilon)\right)$$

for $j\in\{1,2,3\}$ and note that $\lim_{\epsilon\to 0}\left(y_i^{RAT}(q_{i,k}-\epsilon-q_{i,k-1})\middle|\omega_{jk}(\epsilon)\right)=q_{i,k}$; $\lim_{\epsilon\to 0}\left(y_i^{RAT}(q_{i,k+1}-q_{i,k-1}+\epsilon)\middle|\omega_{jk}(\epsilon)\right)=q_{i,k}$; $\lim_{\epsilon\to 0}\left(q_{i,k-1}-\epsilon\middle|\omega_{jk}(\epsilon)\right)=q_{i,k}$; and, because Kastl's (2011) Lemma 1 holds, $\lim_{\epsilon\to 0}\left(\mathbb{P}\left(\omega_{jk}(\epsilon)\right)\right)=0$. 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 $\rho>0$, is the second case of Equation 1.

A.1 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 she demands 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 her bidding costs, $c_i(K_i)$, so these elements do not enter into the bidder's tradeoff, even if she is risk averse. The only random component in her utility function is the price that she must pay for the units that she wins. Simple manipulation as follows shows that Equation 1 simplifies to Equation 2 under these assumptions.

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

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho 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})\\
&\quad=\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=0\\
\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho 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>0
\end{cases}\\
&\implies e^{-\rho\int_0^{q_{i,k}}v^{g,h}(u,\theta_i)du}e^{\rho c_i(K_i)}\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}}\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})\\
&\quad=\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=0\\
e^{-\rho\int_0^{q_{i,k}}v^{g,h}(u,\theta_i)du}e^{\rho c_i(K_i)}\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho\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} & \text{if } \rho>0
\end{cases}\\
&\implies v^{g,h}(q_{i,k},\theta_i)\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}}\middle|b_{i,k}>P^g>b_{i,k+1}\right]\mathbb{P}(b_{i,k}>P^g>b_{i,k+1})\\
&\qquad-\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}}P^g\middle|b_{i,k}>P^g>b_{i,k+1}\right]\mathbb{P}(b_{i,k}>P^g>b_{i,k+1})\\
&\quad=\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=0\\
\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho\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} & \text{if } \rho>0
\end{cases}\\
&\implies\mu_{i,k}=v^{g,h}(q_{i,k},\theta_i)-b_{i,k}\\
&\quad=\frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho 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 P^g q_{i,k}}\middle|b_{i,k}>P^g>b_{i,k+1}\right]}-b_{i,k}\\
&\qquad+\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=0\\
\Omega\frac{1}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}}\middle|b_{i,k}>P^g>b_{i,k+1}\right]}\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho\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} & \text{if } \rho>0
\end{cases}
\end{aligned}
$$

where $\Omega=\frac{1}{\mathbb{P}(b_{i,k}>P^g>b_{i,k+1})}$, which corresponds to Equation 2.

A.2 Market power component

The following shows that Equation 3 holds. The first line follows because $\frac{\partial P^g}{\partial q_{i,k}} = 0$ for $P^g = b_{i,k}$,

$$
\begin{aligned}
& q_{i,k}\Omega \frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \frac{\partial P^g}{\partial q_{i,k}} \mathbb{I}\left(b_{i,k} \geq P^g \geq b_{i,k+1}\right)\right]}{\mathbb{E}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} \\
=& q_{i,k}\Omega \frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \frac{\partial P^g}{\partial q_{i,k}} \mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right]}{\mathbb{E}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} \\
=& q_{i,k}\Omega \frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \frac{\partial P^g}{\partial q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right] \mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)}{\mathbb{E}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} \\
=& q_{i,k}\Omega \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{\partial P^g}{\partial q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right] \mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right) \\
& + q_{i,k}\Omega \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{\partial P^g}{\partial q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right] \mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right) \left(\frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right]}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} - 1\right) \\
& + \frac{q_{i,k}\Omega Cov\left(e^{\rho P^g q_{i,k}}, \frac{\partial P^g}{\partial q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right) \mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} \\
=& q_{i,k}\Omega \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{\partial P^g}{\partial q_{i,k}} \mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right] \\
& + q_{i,k}\Omega \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{\partial P^g}{\partial q_{i,k}} \mathbb{I}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)\right] \left(\frac{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right]}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]} - 1\right) \\
& + \frac{q_{i,k}\Omega Cov\left(e^{\rho P^g q_{i,k}}, \frac{\partial P^g}{\partial q_{i,k}} \middle| b_{i,k} > P^g \geq b_{i,k+1}\right) \mathbb{P}\left(b_{i,k} > P^g \geq b_{i,k+1}\right)}{\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{\rho P^g q_{i,k}} \middle| b_{i,k} > P^g > b_{i,k+1}\right]}
\end{aligned}
$$

which corresponds to Equation 3.

A.3 Allowing for ties

Under Assumptions 4.1 and 4.1, Equation 1 simplifies as follows. First note that Assumption 4.1 implies that $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\frac{\partial y_i}{\partial q_{i,k}} \middle| P^g = b_{i,k+1} \wedge Tie^g\right] = (1 - r_{i,k})$. 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, we have

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du}e^{\rho P^g q_{i,k}}e^{-\rho c_i(K_i)}\left(v^{g,h}(q_{i,k},\theta_i)-P^g\right)\Big|b_{i,k}>P^g>b_{i,k+1}\right]\mathbb{P}(b_{i,k}>P^g>b_{i,k+1})\\
&+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[e^{-\rho\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du}e^{\rho\int_{(q_{i,k-1}+r_{i,k}(q_{i,k}-q_{i,k-1}))}^{q_{i,k}} v^{g,h}(u,\theta_i)du}e^{\rho b_{i,k}\left(q_{i,k-1}+r_{i,k}\left(q_{i,k}-q_{i,k-1}\right)\right)}e^{-\rho c_i(K_i)}\\
&\qquad\left(v^{g,h}\left(q_{i,k-1}+r_{i,k}\left(q_{i,k}-q_{i,k-1}\right),\theta_i\right)-b_{i,k}\right)r_{i,k}\Bigg]\mathbb{P}(P^g=b_{i,k}\wedge Tie^g)\\
&+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[e^{-\rho\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du}e^{-\rho\int_{q_{i,k}}^{(q_{i,k}+r_{i,k+1}(q_{i,k+1}-q_{i,k}))} v^{g,h}(u,\theta_i)du}e^{\rho b_{i,k+1}\left(q_{i,k}+r_{i,k+1}\left(q_{i,k+1}-q_{i,k}\right)\right)}e^{-\rho c_i(K_i)}\\
&\qquad\left(v^{g,h}\left(q_{i,k}+r_{i,k+1}\left(q_{i,k+1}-q_{i,k}\right),\theta_i\right)-b_{i,k+1}\right)\left(1-r^g_{i,k+1}\right)\Bigg]\mathbb{P}(P^g=b_{i,k+1}\wedge Tie^g)\\
=&\begin{cases}
q_{i,k}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\tilde{P}^g(\epsilon)\mathbb{I}\left(b_{i,k}\geq\tilde{P}^g(\epsilon)\geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho=0\\
\frac{1}{\rho}\frac{\partial}{\partial\epsilon}\left(\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[e^{-\rho\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du}e^{\rho\tilde{P}^g(\epsilon)q_{i,k}}e^{-\rho c_i(K_i)}\mathbb{I}\left(b_{i,k}\geq\tilde{P}^g(\epsilon)\geq b_{i,k+1}\right)\right]\right)\Big|_{\epsilon=0} & \text{if } \rho>0
\end{cases}
\end{aligned}
$$

Then, cancelling the deterministic components of the bidder's expected utility,

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

Finally, Assumption 4.1 on the shape of the marginal valuation function implies

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

which corresponds to Equation 5.

B Robustness

Table 4 shows the test results for each pair of models, with the bid shading in Model $S(\rho)$ estimated under the assumption that bidders do not tie, i.e. estimated from Equation 2, and the testing specification identical to the main analysis. Figure 4 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.

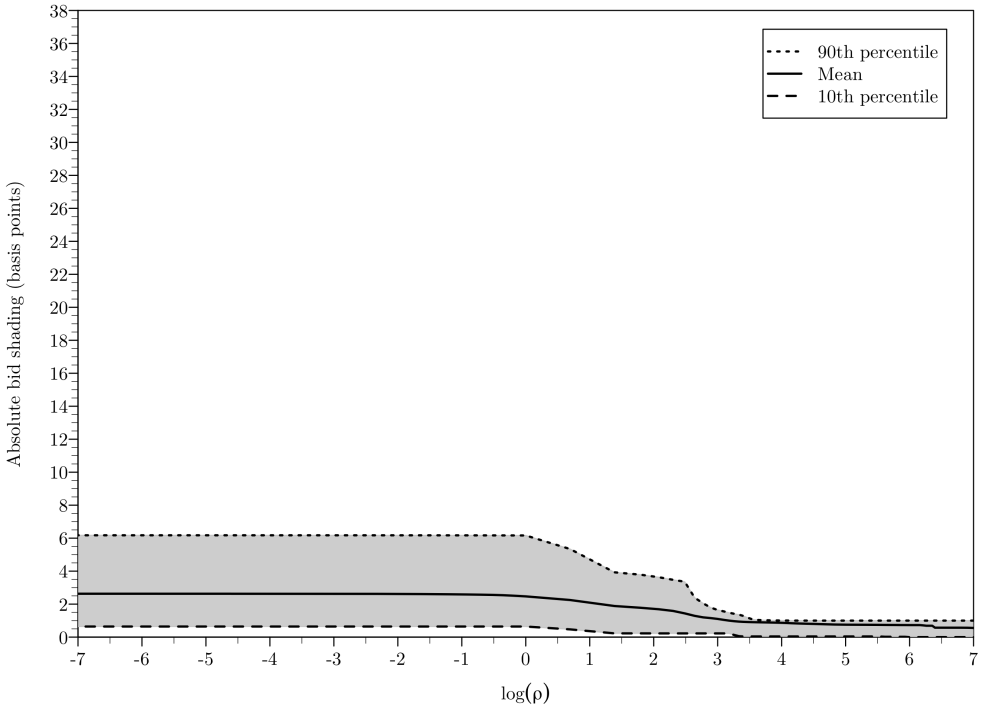
Table 4: Testing results assuming no ties

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

Each entry shows the test statistic for the row being Model 1 and the column being Model 2, so that a negative entry indicates that the row model has better model fit than the column model (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%.

Figure 4: Absolute value of bid shading in Model $S(\rho)$ for $log(\rho) \in [-10, 10]$, June 2010 – May 2012, assuming no ties (observations are weighted by the quantities demanded)

(a) Good A



(b) Good B

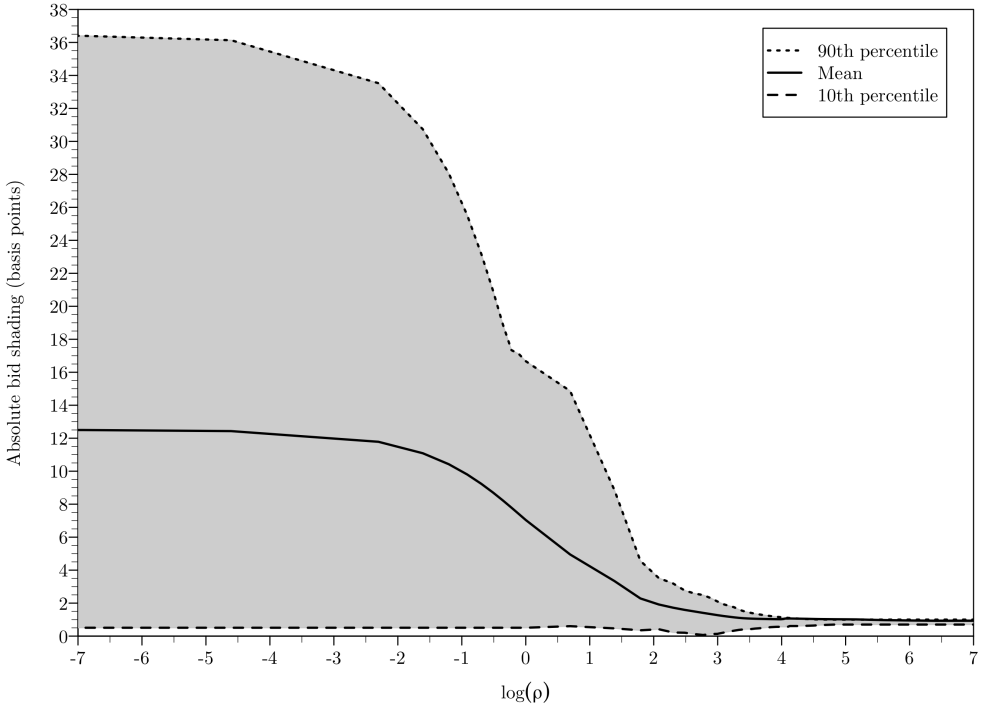


Table 5 shows the test results for each pair of models, with the bid shading in Model $S(\rho)$ estimated under the assumption of the main model and the testing specification identical to the main analysis except that the dummy variable for a last step in a bidder's bid function is excluded from the set of characteristics, $\boldsymbol{x}$.

Table 5: 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

Each entry shows the test statistic for the row being Model 1 and the column being Model 2, so that a negative entry indicates that the row model has better model fit than the column model (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%.

CHAPTER 2
An alternative approach to evaluate bidding in the Bank of England's liquidity auctions

Abstract

This chapter provides an alternative approach to analyse the conventional strategic model of bidding, described in Kastl (2011), in the Bank of England's liquidity auctions. As a measure of model fit, I use the size of the costs of submitting additional bids that is required to explain the data under the assumptions of the model. To estimate the bidding costs, I derive novel necessary conditions that rule out profitable deviations in the bid prices submitted by bidders and use them to estimate bidders' marginal valuation functions. I find that the estimated bidding costs for bidders are relatively large because bidders submit so few bids per auction. However, the physical costs of submitting bids are trivial in the liquidity auctions. This suggests that, under the interpretation of bidding costs as physical costs, the model does not fit this dimension of behaviour well.

This chapter uses the same data and model as Chapter 1. I refer the reader to Chapter 1 for the relevant descriptions.

1 Introduction

In the previous chapter, I evaluate the relative performance of alternative models of behaviour, including Model $S(0)$, a conventional model in which bidders are both strategic and risk neutral. This comparative approach is most useful because one challenge for practitioners is to select the best model among a set of candidates, all of which are approximations to actual behaviour. However, the approach gives no insight to absolute model fit so this chapter analyses an external measure of fit of Model $S(0)$. The model assumes that bidders face "bidding costs", unobserved by the researcher, which deter bidders from making a large number of bids and which explain the few number of bids submitted in practice. The measure of fit that I focus on is the size of these bidding costs required to explain the data under the assumptions of the model. In the BoE's ILTR auctions, the physical costs of submitting bids are trivial.[1] Under this interpretation, I find that the estimated bidding costs for bidders for good B are relatively large, which suggests evidence against this model.

To estimate the bidding costs, I make use of the fact that the number of bids submitted is a strategic choice for a bidder. If she submits K bids, the bidding cost of the $(K+1)$th bid must be larger than the marginal benefit of "fine-tuning" her bid function, which is equal to the difference in her expected utility from submitting what would be the optimal $(K+1)$-step bid function and what would be the optimal K-step one. Having estimated a bidder's full marginal valuation function, we can estimate this marginal benefit of fine-tuning, which provides a lower bound on the bidding cost of the $(K+1)$th bid.

This exercise requires estimates of bidders' full marginal valuation functions, but the condition which rules out profitable deviations in the quantities that bidders demand (Equation 5 in Chapter 1) only identifies bidders' marginal values at the quantities demanded at the steps of their bid functions. I therefore provide a set of novel additional necessary conditions which rule out profitable unilateral deviations in the bid prices submitted by a bidder, holding the rest of her strategy constant. Kastl (2011) derives similar conditions for a deviation in bid price, under the conditions that the support of the auction price is continuous and bidders do not tie on the margin. I instead derive conditions for the case in which the set of possible auction prices is discrete and bidders may ties on the margin, both of which

[1] The following section discusses alternative interpretations in terms of cognitive costs.

characterise my context. These additional conditions set identify the full marginal valuation function. (It is not possible to point identify the full marginal valuation function without further assumptions.) I combine these with the necessary conditions which point identify the marginal valuation function at the quantities demanded at the steps of the bid function (Equation 5 with $\rho = 0$), and use a statistical interpolant to point estimate bidders' full marginal valuation functions in a way which is consistent with these conditions.

I then estimate the expected marginal benefit of submitting an additional bid by the difference in a bidder's expected utility from submitting the optimal $(K+1)$-step bid function and from submitting the optimal K-step bid function. This provides a lower bound on the bidding costs implied by the model.[2] Because the estimation procedure is computationally intensive, I obtain results for bidders who submit one or two bids per auction, which covers 81% of observations.

For a range of marginal valuation functions consistent with the necessary conditions for equilibrium, I find that the average lower bound on the bidding costs for bidders who submit one bid per auction for funds against Level B collateral (67% of observations for B), ranges from £1,800 to £12,000, equivalent to around 3% of average bidder surplus, and is significantly different from zero. I also find that the estimated lower bounds are of similar magnitude when using the same estimation method for the marginal valuation functions as Kastl (2011). These bidding costs seem large when interpreted as the physical costs of submitting bids. These seemingly large costs are explained by the fact that bidders make very few bids in the BoE's liquidity auctions but could increase their surplus by submitting additional bids. The results suggest that the conventional strategic model does not characterise this aspect of bidding behaviour well, at least for the 67% of bidders for good B who submit only one bid. This complements the findings in Chapter 1 that the simpler model of truthful bidding is a better approximation.

In contrast, the results show that the average lower bound on bidding costs for bidders who submit one bid per auction for funds against Level A collateral (68% of observations for A), the estimated lower bounds on bidding costs are very small, equal to approximately £100, equivalent to around 0.2% of average bidder surplus. This is unsurprising given the auction price for these funds is stable across auctions so there are limited gains from further fine-tuning one's bid function; it is also consistent with the findings in Chapter 1 that we cannot discriminate between alternative models of behaviour for bidders for A. Similarly,

[2] An analogous method (which is not relevant to the external measure of fit) provides an upper bound on the bidding costs, equal to the difference in a bidder's expected utility from submitting the optimal K-step bid function and from submitting the optimal $(K-1)$-step bid function.

the average lower bound on bidding costs for bidders who submit two bids per auction for funds is very small for both goods. This is consistent with Kastl's (2012) finding that even a small cost of submitting additional bids can make it optimal to submit few bids.

This chapter proceeds as follows. Section 2 first discusses possible interpretations of the bidding costs. Section 3 provides necessary conditions for equilibrium that, together with Equation 5 in Chapter 1, set identify bidders' marginal valuation functions. Estimating the marginal valuation functions and bidding costs is described in Sections 4 and 5. Section 6 shows the results and Section 7 concludes.

2 Interpretation of bidding costs

Evaluating Model $S(0)$ based on the size of the bidding costs requires an interpretation of these costs.

The most natural interpretation is that they represent the costs of physically submitting the bids. In the ILTR auctions, as in many other settings, bids are submitted through an online platform, called BTender, which is used by bidders to participate in many other BoE operations. So the financial and time costs of submitting bids are trivial.[3]

Alternatively, the bidding costs could be interpreted cognitively. For example, if we assume bidders have infinite experience in participating in the auctions, then they know the average expected utility from submitting a given number of steps based on prior experience, and decide how many steps to submit given the bidding cost of each. Given her choice of number of steps, K_i, a bidder then determines her optimal bids. Under this interpretation, the bidding cost may be interpreted as the cost of calculating the optimal K_i-step bid function. However, this requires a long history of a stable environment, which is not directly applicable to the ILTR setting.

An alternative cognitive interpretation would be that bidders have unlimited cognitive ability, calculate the maximum expected surplus that can be obtained for every possible number of submitted bids, and then choose the optimal number to submit. This would imply that only negligible bidding costs could rationalise the few bids per bidder that are seen in the data. Both cognitive interpretations seem extreme if taken literally, and it is unclear which is preferable.[4]

[3]There is a fixed cost of registering to participate in the ILTR auctions. The bidder must register to become a participant of the BoE's Sterling Monetary Framework and Open Market Operations and typically pledge collateral prior to the auction. These costs are sunk at the time of bidding.

[4]Because bidders can submit any number of bids in the ILTR auctions, we cannot interpret these bidding

3 Set identification of the marginal valuation function

In equilibrium of Model $S(0)$, the observed number of bids submitted by a bidder is optimal, chosen to trade off the marginal cost of an additional bid with the marginal benefit of fine-tuning the bidder's bid function, i.e. the additional expected utility net of bidding costs from submitting an additional bid.

For type θ_i of bidder i, the marginal benefit is equal to the difference in expected utility (net of bidding costs) between submitting the $(K_i + 1)$-step bid function which maximises expected utility conditional on submitting (K_i+1) bids, and submitting the observed K_i-step bid function, which is optimal conditional on submitting K_i bids. This difference therefore provides an estimate of the lower bound of bidder i's bidding cost of an additional bid, $c_i(K_i + 1) - c_i(K_i)$.

To estimate this difference in expected utility, we must first estimate the bidder's unobserved marginal valuation function.

While point identification in a single-unit auction is feasible under certain conditions, identification in multi-unit auctions is more challenging (McAdams, 2008). Equation 5 point identifies a bidder's marginal values at the quantities she demands at the steps of her bid function by ruling out profitable unilateral deviations in the quantity she demands at each step, holding the rest of her strategy constant.

I derive a set of novel necessary conditions that rule out profitable local deviations in the *bid prices* that a bidder submits, in order to set identify the marginal valuation function between the quantities a bidder demands at the steps of her bid function. I also bound the bidder's marginal value for the first unit (i.e. the intercept of the marginal valuation function) from below.

Equilibrium in Model $S(0)$ requires that a bidder cannot profit from a unilateral deviation in the bid price that she submits at a particular step, holding the rest of her strategy fixed. I consider single basis point deviations, above and below the original bid price, as these are the smallest permitted deviations in price that a bidder may make. Equilibrium requires

costs as a constraint imposed by the auction rules. (In this interpretation, the cost of submitting bids of number greater than $\bar{K}$ would be infinite: $c_i(K_i) \to \infty \; \forall K_i > \bar{K}$.) Upper bounds on the number of bids are observed in other settings, but they are never binding as far as I know. For example, in Kastl's (2011) Czech Treasury auctions, the upper bound is 10 and the maximum number of bids by an individual bidder is 9. Similarly, in Hortaçsu and Puller's (2008) study of Texan electricity auctions, "only one firm ever used the maximum number of steps, and that only occurred once for that firm" (McAdams, 2008, citing private communication with Steven Puller).

that the bidder's expected utility given each of these unilateral local deviations is weakly less than her expected utility given her original strategy. The change in expected utility depends on whether the bidder has market power, and the amount of rationing the bidder expects to face if the auction price is set equal to her bid.

To build intuition, I first consider two examples in which a bidder submits a single bid. In both examples, the bidder is a price taker, so her deviation does not affect the distribution of the auction price. In the first example, there is no rationing, so all bids weakly above the auction price are fully allocated, whereas the second example allows for rationing.

In Proposition 1, I then state the pair of conditions that are necessary for equilibrium in Model $S(0)$, allowing for multiple bids and market power. This pair of conditions provides set identification of each bidder's marginal valuation function at all quantities (not just at the quantities she demands at the steps of her bid function). To identify the marginal valuation functions in Model $S(0)$, I assume that the rationing coefficient that bidder i expects to face at a particular price, conditional on the strategy profile, is deterministic. To my knowledge, the constraints imposed by Proposition 2 have not yet been applied in the literature.

By ruling out a profitable unilateral downward deviation in bid price by 1 basis point, Corollary 1 provides set identification of each bidder's marginal value for the first unit (i.e. intercept of their marginal valuation function). This also makes the identifying assumption of a deterministic rationing coefficient, conditional on the strategy profile. I am unaware of any existing studies that have recognised this bound.

3.1 Price-taking bidder, no rationing

I first consider a bidder who submits a single bid on good $g \in \{A, B\}$ at bid price b for a quantity q. In this example, the bidder is a price taker (so she takes the distribution of the auction price as fixed) and there is no rationing (so all bids weakly above the auction price are fully allocated).

I consider two unilateral deviations in bid price to $b - 1$ and $b + 1$, holding the quantity demanded by the bidder (and the number of bids, equal to one) constant.

Equilibrium requires that the bidder cannot profit from either of these unilateral deviations. That is, her expected utility if she follows her original strategy must weakly exceed her expected utility if she instead deviates.

Since she is a price taker and there is no rationing, the downward deviation only changes

her allocation in states in which the auction price is b: she wins the entire quantity that she demands under her original strategy, but wins nothing if she deviates. Equilibrium requires that she weakly prefers to win rather than lose in this case, and so requires that her average marginal value for those units is greater than the price paid, b. The quantity that she wins at all other possible auction prices is unchanged.

Analogously, equilibrium requires that the bidder would prefer to lose at price $b+1$ rather than win the quantity q that she originally demands at price b. She would win at this price given the upward deviation but lose under her original strategy. So equilibrium requires that her average marginal value for those units is less than the price that would be paid, $b+1$.

Combining these two conditions, the bidder's average marginal value for the units up to the quantity that she demands is bounded by her bid price and the integer above her bid price:

$$b \leq \underbrace{\frac{1}{q}\int_0^q v^{g,h}(u,\theta_i)du}_{\text{Average marginal value for quantity demanded}} \leq b+1 \tag{1}$$

This constraint is illustrated in Figure 1 in an example in which the bidder's marginal valuation function is linear. Equilibrium requires that Area G $\geq$ Area H in panel (a) and Area I $\leq$ Area J in panel (b).

3.2 Price-taking bidder, allowing for rationing

For the same price-taking bidder, I now allow for the possibility that the bidder is rationed if the auction price is set equal to her bid. Suppose that the rationing coefficient is deterministic at each auction price, conditional on the strategy profile: $r^g_{i,b}$ is the rationing coefficient at auction price $P^g = b$, conditional on the original strategy profile.[5]

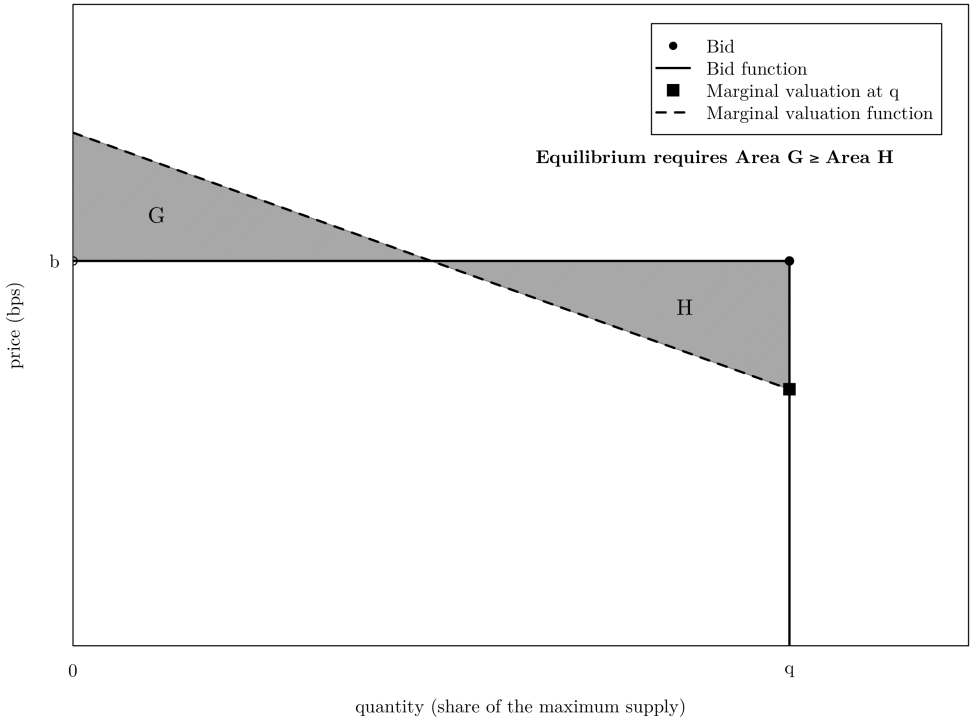
For this bidder, I consider the same two unilateral deviations in bid price to $b-1$ and $b+1$. I denote the respective rationing coefficients at prices $b-1$ and $b+1$ by $r^g_{i,b-1}$ and $r^g_{i,b+1}$.

First consider the downward deviation to $b-1$. The bidder is a price taker, so the deviation only changes her allocation in states in which the auction price is either $b-1$ or b. In the case that $P^g = b-1$, she is fully allocated the quantity that she demands under her original strategy, but rationed if she deviates. In the case that $P^g = b$, she is rationed under her original strategy, and strictly loses if she deviates. So her deviation weakly *reduces* the

[5]One interpretation is that there is very little variation in the rationing coefficient at each auction price, so the bidder believes it to be approximately constant.

Figure 1: Constraints on the marginal valuation function for a single-step price-taking bidder, without rationing

(a) Downward deviation



(b) Upward deviation

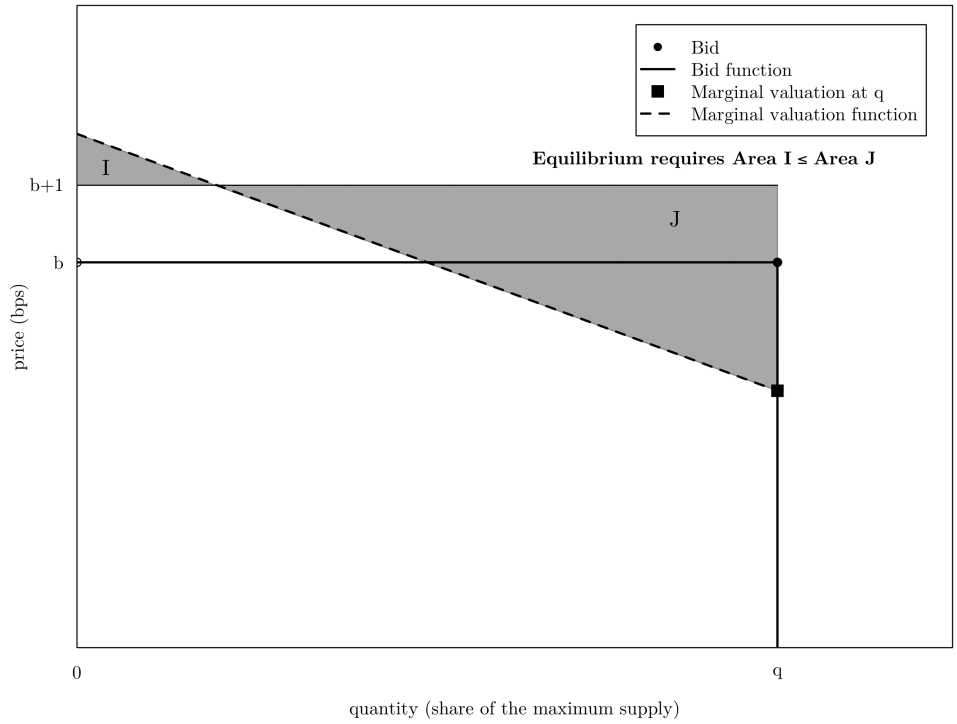
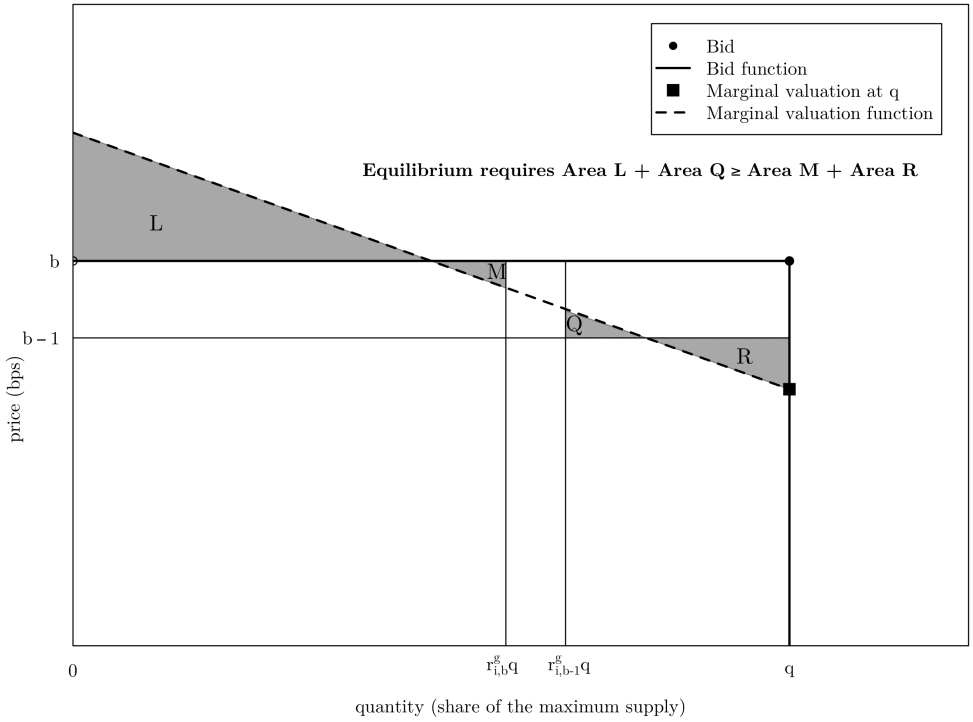
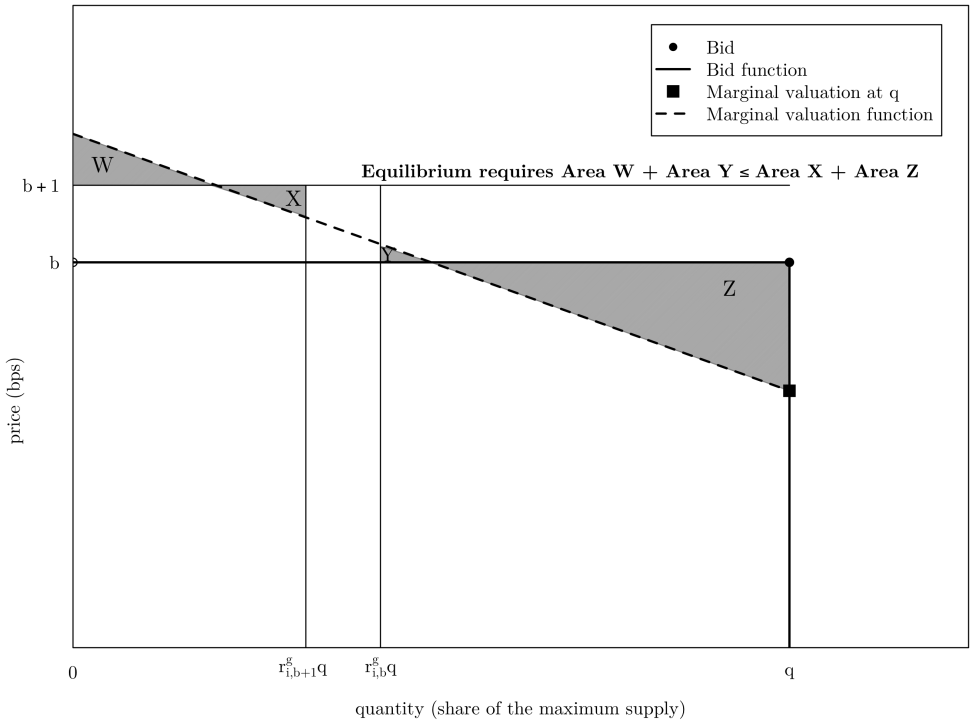


Figure 2: Constraints on the marginal valuation function for a single-step price-taking bidder, with rationing

(a) Downward deviation



(b) Upward deviation



quantity she is allocated at each of the auction prices, $b-1$ and b.

Equilibrium requires that deviating from the bidder's original strategy reduces her expected utility. So her value for the additional units that she is allocated under her original strategy, weighted by the probability of the auction prices occurring, must be weakly greater than the increase in total expected payment that she would make, that is,

$$\underbrace{\mathbb{P}(P^g = b)\int_0^{qr^g_{i,b}} v^{g,h}(u,\theta_i)du + \mathbb{P}(P^g = b-1)\int_{qr^g_{i,b-1}}^{q} v^{g,h}(u,\theta_i)du}_{\text{Expected increase in valuation given demand } q\text{, and bid price } b\text{, relative to } b-1}$$
$$\geq \underbrace{\mathbb{P}(P^g = b)bqr^g_{i,b} + \mathbb{P}(P^g = b-1)(b-1)q\left(1 - r^g_{i,b-1}\right)}_{\text{Increase in expected payment given demand } q\text{, and bid price } b\text{, relative to } b-1} \quad (2)$$

Analogously, equilibrium requires that the bidder cannot profit from a unilateral upward deviation in bid price by one unit to $b+1$. The deviation only changes her allocation in states in which the auction price is either b or $b+1$, weakly increasing it in both cases. Given $P^g = b$, she is rationed under her original strategy but fully allocated if she deviates. Given $P^g = b+1$, she strictly loses under her original strategy but is rationed if she deviates.

Equilibrium requires that the bidder's value for the additional units that she is allocated if she deviates, weighted by the probability of the auction prices occurring, is weakly *lower* than the expected increase in total payment that she would make, that is,

$$\underbrace{\mathbb{P}(P^g = b+1)(b+1)qr^g_{i,b+1} + \mathbb{P}(P^g = b)bq\left(1 - r^g_{i,b}\right)}_{\text{Increase in expected payment given demand } q\text{, and bid price } b+1\text{, relative to } b}$$
$$\geq \underbrace{\mathbb{P}(P^g = b+1)\int_0^{qr^g_{i,b+1}} v^{g,h}(u,\theta_i)du + \mathbb{P}(P^g = b)\int_{qr^g_{i,b}}^{q} v^{g,h}(u,\theta_i)du}_{\text{Expected increase in valuation given demand } q\text{, and bid price } b+1\text{, relative to } b} \quad (3)$$

These two conditions are illustrated in Figure 2, in an example in which the bidder's marginal valuation function is linear and the distribution of the auction price is uniform. Equilibrium requires that Area L + Area Q $\geq$ Area M + Area R in panel (a) and that Area W + Area Y $\leq$ Area X + Area Z in panel (b).

3.3 Equilibrium conditions

For the price-taking bidder considered above, the incentive to deviate depends only on the effect on the quantity that she is allocated at each auction price. Equilibrium requires that the expected change in total payment from a unilateral local deviation in bid price exceeds the expected change in valuation, which deters the bidder from deviating.

Proposition 1 generalises this condition to also account for the fact that each bidder may submit multiple bids, and may have market power so that their deviation affects the distribution of the auction price. I consider two, single basis point deviations: one below the original bid price, and one above. Equations 4 and 5 specify necessary conditions for equilibrium, which rule out that these unilateral deviations are profitable.

In Model $S(0)$, the equilibrium strategy profile is $\boldsymbol{\sigma}(\boldsymbol{\Theta})$. Let $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta})$ denote the strategy profiles which are identical to $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ except that the bid prices submitted by type θ_i of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ at step k, in her K_i-step bid function, are equal to $b_{i,k} - 1$ and $b_{i,k} + 1$, respectively. For ease of notation, I write $\boldsymbol{\sigma}(\boldsymbol{\Theta}|\theta_i) = \boldsymbol{\sigma}(\boldsymbol{\Theta})$, and similarly for $\boldsymbol{\sigma}'(\boldsymbol{\Theta}|\theta_i)$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta}|\theta_i)$.

I make the following identifying assumption that the rationing coefficient is deterministic at each auction price, conditional on the strategy profile.[6] This can be loosely interpreted as bidders believing that the rationing coefficient at each auction price has little variation around its expected value, and so can be treated as approximately deterministic. In Appendix A, I show how this assumption permits identification.

The rationing coefficients for bidder i at step $k \in \{1, ..., K_i\}$, conditional on the strategy profiles $\boldsymbol{\sigma}(\boldsymbol{\Theta}|\theta_i), \boldsymbol{\sigma}'(\boldsymbol{\Theta}|\theta_i)$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta}|\theta_i)$ are deterministic and are denoted $r_{i,k}, r'_{i,k}$, and $r''_{i,k}$, respectively.

Proposition 1 follows from this assumption and the assumptions of Model $S(0)$.

Proposition 1 (Necessary condition on bid price deviations in Model $S(0)$) *In Model $S(0)$ under Assumption 3.3, in any type-symmetric Bayesian Nash Equilibrium, for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, every step $k \in \{1, ..., K_i\}$ in the K_i-step*

[6] The necessary conditions for equilibrium without the identifying assumptions are Equations 12 and 15 in Appendix A.

bid function in $\sigma_i(\theta_i)$ *where* $b_{i,k}-1 > b_{i,k+1}$ *and* $b_{i,k}+1 < b^g_{i,k-1}$*, must satisfy*[7]

$$\begin{aligned}
&\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u,\theta_i)du - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}\right)\right)\mathbb{P}\left(P^g = b_{i,k}\middle|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\\
&\quad+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - (b_{i,k}-1)q_{i,k}\right)\mathbb{P}\left(P^g = b_{i,k}-1\middle|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\\
&\geq\left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k-1}\right)\mathbb{P}\left(P^g = b_{i,k}\middle|\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)\\
&\quad+\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r'_{i,k}} v^{g,h}(u,\theta_i)du\right.\\
&\qquad\left.-(b_{i,k}-1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r'_{i,k}\right)\right)\mathbb{P}\left(P^g = b_{i,k}-1\middle|\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \qquad (4)
\end{aligned}$$

and

$$\begin{aligned}
&\left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - (b_{i,k}+1)q_{i,k-1}\right)\mathbb{P}\left(P^g = b_{i,k}+1\middle|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\\
&\quad+\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u,\theta_i)du\right.\\
&\qquad\left.-b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}\right)\right)\mathbb{P}\left(P^g = b_{i,k}\middle|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\\
&\geq\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r''_{i,k}} v^{g,h}(u,\theta_i)du\right.\\
&\qquad\left.-(b_{i,k}+1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r''_{i,k}\right)\right)\mathbb{P}\left(P^g = b_{i,k}+1\middle|\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)\\
&\quad+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k}\right)\mathbb{P}\left(P^g = b_{i,k}\middle|\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right) \qquad (5)
\end{aligned}$$

Proof. See Appendix A. ■

As discussed in Section 4.2, each of the probabilities and rationing coefficients in Equations 4 and 5 can be estimated and $b_{i,k}, q_{i,k}$ and $q_{i,k-1}$ are observed. Approximating the integrals with summations over an arbitrarily fine quantity grid, these two equations therefore impose linear constraints on the marginal valuation functions of bidders that rationalise observed behaviour in Model $S(0)$.

[7]The conditions for bidders who submit adjacent bid prices are analogous, and shown in Appendix A.3.

By ruling out profitable unilateral deviations in bid price, it is also possible to bound each bidder's marginal value for the first unit (i.e. intercept of their marginal valuation function) from below.

Corollary 1 (Necessary condition on marginal valuation function intercepts in Model $S(0)$) *In Model $S(0)$ under Assumption 3.3, in any type-symmetric Bayesian Nash Equilibrium, for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, step $k = 1$ in the K_i-step bid function in $\sigma_i(\theta_i)$ must satisfy*

$$v^{g,h}(0, \theta_i) \geq b^g_{i,1} - 1 \tag{6}$$

under both the assumptions that $\mathbb{P}\left(P^g = b^g_{i,1} \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) > 0$ or $\mathbb{P}\left(P^g = b^g_{i,1} - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) > 0$ and that $r_{i,1} \neq 0$.

Proof. See Appendix B. ■

Equation 6 places a further restriction on the set of marginal valuation functions that rationalise observed behaviour. This constraint is slack for a bidder who submits bids at prices below her marginal value at the quantity she demands at the first step of her bid function because the marginal valuation function is downward sloping (Assumption 3.1.1), i.e. for whom $\mu^g_{i,1} \geq 0$ (which can be determined by Equation 5). However, for bidders who submit bids at prices above their marginal value at the quantity they demand at the first step of their bid functions, i.e. $\mu^g_{i,1} < 0$, Corollary 1 restricts the set of marginal valuation functions consistent with equilibrium in Model $S(0)$.

4 Estimating marginal valuation functions

Equations 5 and 4-6 describe necessary relationships between a bidder's observed bids and her unobserved marginal valuation function in equilibrium of Model $S(0)$, for a particular realisation of her signal. As described in Section 4.2, I use the first equation to obtain point estimates of each bidder's marginal values at the quantities they demand at the steps of their bid function. The points of a bidder's marginal valuation function between the quantities she demands at the steps cannot be point identified, but are set identified. I therefore interpolate the remaining parts of the bidder's marginal valuation function in a way which is consistent with Proposition 1 and Corollary 1.

There is a range of possible marginal valuation functions that satisfy Equations 4-6: the intercept for bidder i could be any value greater than $(b^g_{i,1} - 1)$ and, for each of these

intercepts, there is potentially a large range of marginal valuation functions consistent with equilibrium in Model $S(0)$.

If the point estimate of the bidder's marginal value at the quantity she demands at the first step of her bid function is greater than her bid, i.e. $\mu^g_{i,1} \geq 0$, I set the intercept equal to this point estimate. For all other observations, the lower bound specified in Equation 6 is binding so I produce three estimates corresponding to three different intercepts: the price of the first step of the bidder's bid function plus a "wedge", $w \in \{1, 5, 10\}$.[8]

$$\hat{v}_i(0) = \max\{b_{i,1} + w, \hat{v}_i(q_{i,1})\} \text{ for } w = \{1, 5, 10\} \tag{7}$$

where $\hat{v}_i(q)$ is the estimated marginal value of bidder i at quantity q.

For each of these intercepts, I estimate a marginal valuation function that passes through the point estimates at the quantities demanded at the steps of the bid function, using Gaussian Process Regression (GPR). This is a Bayesian non-parametric statistical interpolant that assumes a Gaussian prior, updated using the point estimates of the marginal values and the specified intercept (Rasmussen and Williams, 2006).[9] It is more suitable than other methods such as linear regression as there is no prima facie measurement error in the point estimates. It also provides more flexibility as there is no strong reason to believe that the marginal valuation function is linear and the prior is defined over the entire function, rather than the parameters. Intuitively, it chooses the most likely function that passes through the specified points.[10]

I then use minimum-distance estimation to find the marginal valuation function nearest to the GPR estimate that is consistent with the monotonicity assumption (Assumption 3.1.1) and Proposition 1; I denote the estimator by MD-GPR.[11]

[8]The intercept of the marginal valuation function is not point identified, and there is no appropriate loss function to guide the choice of intercept within the set specified in Corollary 1. I therefore use three specifications, with the values of w motivated by the standard deviation of bids submitted in the ILTR auctions.

[9]For an arbitrarily fine grid of quantities given by $\boldsymbol{q} = (q_1, q_2, \ldots, q_N)$, where $q_1 < q_2 < ... < q_N, q_1 = 0, q_N \geq q_{i,K_i}$, and $q_{i,k} \in \{q_1, ..., q_N\}\ \forall k \in \{1, ..., K_i\}$, the GPR prior probability density of the vector of function values $\boldsymbol{v_i}(\boldsymbol{q}) = (v_i(q_1), v_i(q_2), ..., v_i(q_N))$ is jointly Gaussian, with a covariance matrix defined by a kernel function. The smaller is the distance $|q - q_{i,k}|$, the greater is the covariance between the predicted marginal value at quantity q and the marginal value point estimate, $\hat{v}_i(q_{i,k}), k = \{1, \ldots, K_i\}$.

[10]Alternative GPR specifications, including different kernel functions for the covariance matrix that determine the smoothness of the function, make little difference to the results.

[11]This finds the estimate, nearest to the GPR estimate, that is consistent with the inequality constraints specified in Equations 4 and 5, approximating the integrals sums over quantity units of 0.1% of the maximum supply.

I assume that the bidder's marginal value for quantities larger than the total quantity that she demands equals zero, that is, $\hat{v}_i(q, \theta_i) = 0 \, \forall \, q > q_{i,K_i}$. This reduces the expected utility from submitting an additional bid and so means the results are conservative estimates of the lower bounds of the bidding costs implied by Model $S(0)$.

To compare my results to Kastl (2011), I also estimate the marginal valuation function by a step function which is an upper envelope of the point estimates of the bidder's marginal values at the quantities demanded at the steps of the bid function.[12] In this specification, the marginal values for all quantities across the first step of the bidder's bid function are set equal to the point estimate of the marginal value at the quantity demanded at the first step. And so, for bidders who bid above their marginal value (by more than 1 basis point) at the first step of their bid function, this estimate violates the conditions specified in Proposition 2 and Corollary 1.[13] Under the identifying assumption that I make of a deterministic rationing coefficient, Kastl's (2011) method underestimates the average steepness of the marginal valuation function across the first step of the bid function.

An example of the four estimated marginal valuation functions—three estimated by MD-GPR and the fourth estimated by the upper envelope which is stepped—is shown for a bidder who submits a bid function with two steps for good B in Figure 3. The first three estimates all satisfy the necessary conditions for equilibrium in Model $S(0)$ described in Propositions 1 and 2.

5 Estimating bidding costs

In Model $S(0)$, a bidder who submits K_i bids has chosen to do so to maximise her expected utility: it must be that the marginal cost of submitting an additional bid, equal to $c_i(K_i + 1) - c_i(K_i)$, is greater than the marginal benefit of fine-tuning, equal to the gain in expected utility (net of bidding costs) from submitting the (K_i+1)-step bid function which maximises expected utility conditional on submitting (K_i+1) bids, rather than the K_i-step bid function that she uses. This condition allows us to estimate a lower bound on the size of the bidding costs implied by Model $S(0)$.

For a bidder who submits K_i bids, the K_i-step bid function which maximises expected utility conditional on submitting K_i bids is her observed bid function under the assumptions

[12] Specifically, this is a step function with the quantities at each step coinciding with those of the bid function. The marginal values at the quantities demanded at the steps of the bid function are point identified by Equation 5 (with $\rho = 0$) and the marginal values for quantities greater than the total quantity demanded by the bidder are equal to zero.

[13] Both Kastl (2011) and I find this to be economically relevant.

of Model $S(0)$. I estimate the bidder's expected utility given her strategy by estimating the distribution of the auction price using the resampling technique described in Section 4.2. I use this distribution to estimate the bidder's expected utility (net of bidding costs) of using a K_i-step bid function for each of the estimates of the marginal valuation functions.

The (K_i+1)-step bid function which maximises her expected utility conditional on submitting $(K_i + 1)$ bids is unobserved. I implement a grid search across all bid prices and quantity increments of 0.25% to estimate this optimal $(K_i + 1)$-step bid function. For each of the $(K_i + 1)$-step bid functions on this grid, and for each estimate of the bidder's marginal valuation function, I estimate the bidder's beliefs over the distribution of the auction price conditional on the bid function, and then use this distribution to estimate the bidder's expected utility (net of bidding costs). The estimate of the optimal (K_i+1)-step bid function is the one which maximises the bidder's expected utility.

The difference in the estimated expected utilities from the (K_i+1)- and K_i-step optimal bid functions provides an estimate of the lower bound on bidding costs for each estimate of the bidder's marginal valuation function.[14]

To estimate the standard errors, I follow Backus, Conlon and Sinkinson (2021) and Roussille and Scuderi (2022) by treating the optimal bid function for each number of steps as known.

6 Results

Table 1 shows the estimated lower bounds on the bidding costs of submitting additional bids for goods A and B.[15] These lower bounds are estimated for the three MD-GPR estimates of the marginal valuation function: $\max\left\{b^g_{i,1} + w, \hat{v}^g_i(q^g_{i,1})\right\}$ for $w = \{1, 5, 10\}$, and for the step function which is the upper envelope of the point estimates of the bidder's marginal values at the quantities she demands at the steps of her bid function.

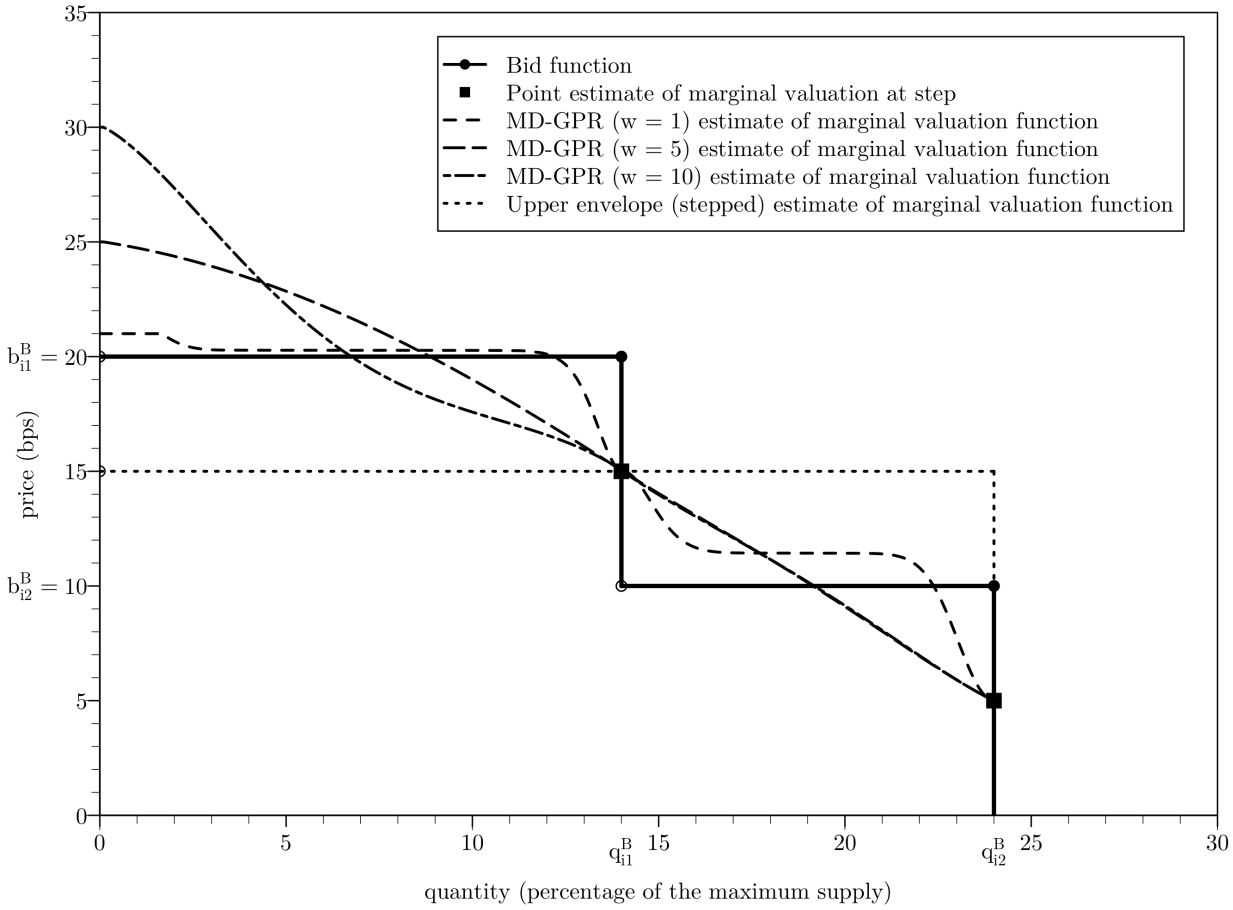
The estimated bidding costs vary across goods, numbers of steps, and intercepts, and within these subgroups (shown by the large standard deviations). This variation can be interpreted in terms of the variation in the expected benefit of fine-tuning one's bid function from submitting additional bids.[16]

[14] We can also estimate an upper bound on the bidding cost of the K_ith bid, by comparing the estimated expected utilities (net of bidding costs) of submitting the optimal K_i- and $(K_i + 1)$-step bid functions.

[15] The lower bounds for the second step are estimated for every bidder, for every auction in which they participate and submit a bid function with one step, in June 2010 – May 2012. The lower bounds for the third step are estimated, analogously, for every bidder, for every auction in which they participate and submit a bid function with two steps.

[16] The heterogeneity itself if difficult to explain if the costs are interpreted as the physical costs of sub-

Figure 3: Example of four estimated marginal valuation functions of a bidder, with using MD-GPR and an upper envelope.



The results show that the average lower bound on the bidding cost of submitting a second step for good B is large: bidders who submit one bid for good B are giving up between £1,800 and £12,707 on average by only submitting one bid. If the costs are interpreted as physical costs, which are trivial in practice, this suggests that Model $S(0)$, which assumes both that bidders are strategic and that they are risk neutral, does not provide a good characterisation of the behaviour of these bidders. This is consistent with the results in Chapter 1: for bidders for good B, a "truthful bidding" model in which bidders' bids correspond to their true marginal values for liquidity fit the data better than Model $S(0)$.

In contrast, the average lower bound on bidding costs for bidders for good A are very small, even for bidders who submit one bid. This can be explained by the fact that there are limited benefits for these bidders from submitting additional bids, because the range of possible clearing prices is so narrow. The average clearing price for A in June 2010 – May 2012 is 1.38bps and its standard deviation is 1.82bps in the 3-month term auctions (and these statistics are respectively 0.38bps and 0.74bps in the 6-month term auctions), so bidders gain little from fine-tuning.

Finally, the bidding cost of submitting a third step is also small. This suggests that bidders may only forego a small amount of surplus by submitting a two-step bid function as an approximation to what would be the optimal bid function in the absence of bidding costs. This follows similar intuition to Wilson's (1993) finding that a monopolist can approximate an optimal non-linear multi-part tariff with four or five two-part tariffs.

The findings for good B, including the results for the marginal valuation functions estimated as step functions, differ from those found by Kastl (2011), in a study of Czech Treasury auctions. In that context, he estimates the lower bound on the bidding costs of submitting a second step to be between $2 and $150. There are some clear institutional differences which might explain why Model $S(0)$ seems to better characterise bidding behaviour in the Czech Treasury auctions. For example, bidders appear to face less uncertainty in Kastl's (2011) setting, in which only one good is auctioned and supply uncertainty (generated by non-competitive bids) is limited. However, there are some comparable features, including the average number of bids submitted per bidder.[17]

mitting bids, which are constant across bidders and bids in the ILTR auctions.

[17] In Kastl's (2011) setting, the average is 1.41 for "Small" bidders (defined as demanding less than 5% of the auction) and is 2.95 for "Large" bidders (the remaining set).

Table 1: Estimated lower bound on bidding costs implied by Model $S(0)$, £, June 2010 – May 2012

Lower bound of cost of submitting 2nd step (£)				
	Good A		Good B	
Estimation type	Mean	Std. Dev.	Mean	Std. Dev.
MD-GPR ($w = 1$)	103.7	206.0	9557.6	41734.4
	(24.0)	(39.8)	(4776.7)	(13107.6)
MD-GPR ($w = 5$)	183.6	365.5	1976.7	5884.7
	(41.9)	(71.2)	(700.6)	(1423.8)
MD-GPR ($w = 10$)	129.1	260.3	1839.1	5222.0
	(28.4)	(44.3)	(640.4)	(1387.9)
Upper envelope (stepped)	61.9	138.8	12707.5	41325.1
	(19.6)	(34.2)	(6680.2)	(13797.1)
Lower bound of cost of submitting 3rd step (£)				
	Good A		Good B	
Estimation type	Mean	Std. Dev.	Mean	Std. Dev.
MD-GPR ($w = 1$)	0.9	2.5	42.3	72.1
	(0.6)	(0.9)	(34.6)	(33.8)
MD-GPR ($w = 5$)	0.5	2.1	106.8	168.7
	(0.5)	(1.3)	(81.5)	(75.6)
MD-GPR ($w = 10$)	0.3	1.4	110.3	134.9
	(0.3)	(0.9)	(64.3)	(47.6)
Upper envelope (stepped)	0.0	0.0	0.0	0.0
	(0.0)	(0.0)	(0.0)	(0.0)

Notes: For each "wedge", $w \in \{1, 5, 10\}$, the intercept of bidder i's marginal valuation function for good $g \in \{A, B\}$ is equal to $\hat{v}_i^g(0) = \max\left\{b_{i,1}^g + w, \hat{v}_i^g(q_{i,1}^g)\right\}$. Standard errors in parentheses.

Table 2: Estimated lower bound on bidding costs implied by Model $S(0)$, percentage of average bidding surplus, June 2010 – May 2012

Lower bound of cost of submitting 2nd step (% of average bidder surplus)				
	Good A		Good B	
Estimation type	Mean	Std. Dev.	Mean	Std. Dev.
MD-GPR ($w = 1$)	0.379	0.547	3.262	6.728
	(0.071)	(0.098)	(1.062)	(1.902)
MD-GPR ($w = 5$)	0.316	0.466	2.336	6.278
	(0.059)	(0.071)	(0.998)	(3.299)
MD-GPR ($w = 10$)	0.106	0.166	2.232	4.816
	(0.021)	(0.029)	(0.776)	(2.342)
Upper envelope (stepped)	0.144	0.326	3.377	8.344
	(0.047)	(0.079)	(1.380)	(3.405)
Lower bound of cost of submitting 3rd step (% of average bidder surplus)				
	Good A		Good B	
Estimation type	Mean	Std. Dev.	Mean	Std. Dev.
MD-GPR ($w = 1$)	0.003	0.009	0.096	0.151
	(0.002)	(0.004)	(0.073)	(0.066)
MD-GPR ($w = 5$)	0.001	0.004	0.237	0.330
	(0.001)	(0.003)	(0.160)	(0.134)
MD-GPR ($w = 10$)	0.002	0.007	0.171	0.234
	(0.002)	(0.004)	(0.113)	(0.097)
Upper envelope (stepped)	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)

Notes: For each "wedge", $w \in \{1, 5, 10\}$, the intercept of bidder i's marginal valuation function for good $g \in \{A, B\}$ is equal to $\hat{v}_i^g(0) = \max\left\{b_{i,1}^g + w, \hat{v}_i^g(q_{i,1}^g)\right\}$. Standard errors in parentheses.

7 Conclusion

In this chapter, I have analysed the conventional strategic model, commonly and increasingly used in the literature, which assumes that bidders best respond to their beliefs about the economic environment and other participants' behaviour and that bidders are risk neutral. To complement the analysis in Chapter 1, which analyses the model's relative model fit, I evaluated an absolute measure of fit of the model in terms of one of its key characteristics: the size of the costs of submitting additional bids, or "bidding costs". These are included in the conventional model to explain the fact that bidders submit very few bids in many multi-unit auction settings, including the Bank of England's liquidity auctions. If the marginal benefits of fine-tuning one's bid function by submitting an additional bid are small, submitting few bids is optimal.

In practice, the physical costs of submitting bids are trivial in the Bank of England's auctions. Under this interpretation, the large bidding costs that I estimate for some bidders suggests that the conventional model of bidding does not describe the behaviour for these bidders well.

Overall, this analysis can be seen as indicative of the lack of fit of Model $S(0)$ to bidding behaviour in the Bank of England's liquidity auctions. Importantly, further work is required to determine an appropriate size for the bidding costs.

Appendices

A Proof of Proposition 1

In Model $S(0)$, the equilibrium strategy profile is $\boldsymbol{\sigma}(\boldsymbol{\Theta})$. Let $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta})$ denote the strategy profiles which are identical to $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ except that the bid prices submitted by type θ_i of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ at step k, in her K_i-step bid function for good g, are equal to $b_{i,k} - 1$ and $b_{i,k} + 1$, respectively. For ease of notation, I write $\boldsymbol{\sigma}(\boldsymbol{\Theta}|\theta_i) = \boldsymbol{\sigma}(\boldsymbol{\Theta})$, and similarly for $\boldsymbol{\sigma}'(\boldsymbol{\Theta}|\theta_i)$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta}|\theta_i)$.

I first derive Equation 4, ruling out a profitable unilateral *downward* deviation in bid price of one unit. I then derive Equation 5, ruling out a profitable analogous *upward* deviation.

A.1 Downward deviation

Equilibrium of Model $S(0)$ requires that the deviation of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ at step k to a bid price of $b_{i,k} - 1$ rather than $b_{i,k}$ weakly reduces her expected utility, and so requires

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)\right] \geq \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i(\boldsymbol{\sigma}'(\boldsymbol{\Theta})|\theta_i)\right] \\
&\implies \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\int_0^{y_i^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u, \theta_i)du - P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))y_i^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) - c_i(K_i)\right] \\
&\quad \geq \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[\int_0^{y_i^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))} v^{g,h}(u, \theta_i)du - P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))y_i^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) - c_i(K_i)\right] \qquad (8)
\end{aligned}$$

By the Law of Iterated Expectations (LIE),

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)\right] \\
&\quad = \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\middle|P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) \notin \{b_{i,k} - 1, b_{i,k}\}\right)\right]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) \notin \{b_{i,k} - 1, b_{i,k}\}\right) \\
&\qquad + \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\middle|P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}\right)\right]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}\right) \\
&\qquad + \mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\middle|P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k} - 1\right)\right]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k} - 1\right) \qquad (9)
\end{aligned}$$

and $\mathbb{E}_{\boldsymbol{\Theta}_{-i}}\left[U_i(\boldsymbol{\sigma}'(\boldsymbol{\Theta})|\theta_i)\right]$ can be decomposed analogously.

First suppose that the bid prices at steps k and $k-1$ are non-adjacent so that $b_{i,k} - 1 > b_{i,k+1}$.

I consider the case in which $b_{i,k} - 1 = b_{i,k+1}$ in Appendix A.3. Using the LIE, the bidder's expected utilities given her strategy and deviation, respectively, can be written as

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)\right] \\
&= \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du \\
&\qquad - P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))y_i(\boldsymbol{\sigma}(\boldsymbol{\Theta}))\Bigg| P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) \notin \{b_{i,k}-1, b_{i,k}\}\Bigg]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) \notin \{b_{i,k}-1, b_{i,k}\}\right) \\
&+ \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du \\
&\qquad - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\right)\Bigg| P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}\Bigg]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}\right) \\
&+ \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - (b_{i,k}-1)q_{i,k}\middle| P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}-1\right]\mathbb{P}\left(P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}-1\right) \\
&- c_i(K_i)
\end{aligned} \tag{10}
$$

and

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}'(\boldsymbol{\Theta})|\theta_i)\right] \\
&= \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{y_i(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du \\
&\qquad - P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))y_i(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))\Bigg| P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) \notin \{b_{i,k}-1, b_{i,k}\}\Bigg]\mathbb{P}\left(P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) \notin \{b_{i,k}-1, b_{i,k}\}\right) \\
&+ \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k-1}\middle| P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k}\right]\mathbb{P}\left(P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k}\right) \\
&+ \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du - (b_{i,k}-1)\Big(q_{i,k-1}+ \\
&\qquad (q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)\Big)\Bigg| P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k}-1\Bigg]\mathbb{P}\left(P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k}-1\right) \\
&- c_i(K_i)
\end{aligned} \tag{11}
$$

where $r^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)$ and $r^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)$ are the rationing coefficients for good g given strategy profiles $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ (and $\Theta_i = \theta_i$); $q_0 = 0$.

First note that the probability of the event $P^g \in \{b_{i,k}-1, b_{i,k}\}$, and the allocation given each

auction price $P^g \notin \{b_{i,k}-1, b_{i,k}\}$ are identical under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$.

To show this, I split the distribution of the auction price of good g into three collectively exhaustive sets, $(P^g < b_{i,k}-1)$, $(b_{i,k} < P^g)$ and $P^g \in \{b_{i,k}-1, b_{i,k}\}$, and consider each in turn. The Product-Mix Auction uses the information from the submitted bids and the BoE's supply curves to find the competitive equilibrium, assuming bids correspond to bidders' marginal values and the supply curves represent the BoE's marginal costs of supply (see Section 2.1). The auction price of good g is the maximum of the highest losing bid on good g and the marginal cost of supply of good g at the quantity allocated (equal to 0 for good A and equal to the sum of P^A and the relative supply of B, evaluated at the quantity allocated of B).

For all signal realisations such that $P^g < b_{i,k} - 1$ under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, bidder i's deviation to a bid price of $b_{i,k} - 1$ at step k does not affect the competitive equilibrium, so the outcome is identical under the two strategy profiles, $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$.

Similarly, for all signal realisations such that $P^g > b_{i,k}$ under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, bidder i's deviation to a bid price of $b_{i,k} - 1$ at step k also does not affect the competitive equilibrium, so the outcome is identical under the two strategy profiles, $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$. Now consider signal realisations such that $P^g = b_{i,k}$ under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$.

First consider $g = A$. The auction price of A is the maximum of the marginal losing bid on A and 0. If $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k} > 0$, then bidder i's bid at step k is marginal under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$. It follows that, under $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$, his bid at step k is either marginal (in which case $P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k} - 1$) or his bid at step k is strictly unallocated (in which case $P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k}$). So if $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}$ then $P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) \in \{b_{i,k} - 1, b_{i,k}\}$.

Now consider $g = B$. There are three subcases:

(i) If bidder i's deviation has no effect on the auction price of good A, i.e. $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))$, the result immediately follows: B's auction price is determined by the intersection of the aggregate bid function for good B and the difference between the relative supply of B and P^A. Bidder i's deviation has no effect on this intersection (if the quantity he demands at step k is less than the amount of unallocated demand at P^B under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$) or his deviation reduces the intersection to $P^g = b_{i,k} - 1$.

(ii) If i's deviation reduces the auction price of good A, i.e. $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) > P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))$, then the allocation of good A is strictly larger under $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ than under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$. (The reason is that P^A is the maximum of the marginal losing bid on A and 0. So a bid on A must have been unallocated under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ but allocated under $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$.) If $P^B(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) < b_{i,k}-1$, then

the quantity allocated of good B under $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ is also weakly larger than the quantity allocated of good B under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, so the total allocation across both goods strictly increases. But $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) > P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))$ implies $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) > 0$, which can only occur if the entire maximum supply is allocated under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$. So we have a contradiction: if $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) > P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))$, it is not possible that $P^B(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))) < b_{i,k} - 1$. Moreover, bidder i's deviation weakly reduces demand at each pair of auction prices (P^A, P^B) so cannot increase the price of good B. Overall, if $P^A(\boldsymbol{\sigma}(\boldsymbol{\Theta})) > P^A(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))$, then $P^B(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) \in \{b_{i,k} - 1, b_{i,k}\}$.

(iii) Bidder i's deviation weakly reduces demand at each pair of auction prices (P^A, P^B), so cannot increase the auction price of good A.

It follows from these three subcases that if $P^B(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k}$ then $P^B(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) \in \{b_{i,k}-1, b_{i,k}\}$.

For analogous reasoning, if $P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) = b_{i,k} - 1$ then $P^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta})) = b_{i,k} - 1$ for $g \in \{A, B\}$. Taken together, the probability of the event $P^g \notin \{b_{i,k} - 1, b_{i,k}\}$, and the allocation of good g given each auction price $P^g \notin \{b_{i,k} - 1, b_{i,k}\}$ are identical under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$.

Equilibrium therefore requires

$$
\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du \\
&\qquad - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\right)\Bigg| P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) = b_{i,k}\Bigg]\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) = b_{i,k}\right) \\
&+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - (b_{i,k}-1)q_{i,k}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) = b_{i,k}-1\right) \\
&\geq \left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k-1}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) = b_{i,k}\right) \\
&+\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}'(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du \\
&\qquad - (b_{i,k}-1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)\right)\Bigg| P^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) = b_{i,k}-1\Bigg]\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) = b_{i,k}-1\right)
\end{aligned}
\tag{12}
$$

Equation 12 cannot be readily estimated as it requires evaluation of an integral with a random interval. I therefore make Assumption 3.3 that the rationing coefficient is deterministic at each auction price, conditional on the strategy profile. Under this assumption, the changes in

the bidder's gross utility, at each auction price, conditional on her strategy, are deterministic, and can be estimated.

Under this assumption, equilibrium requires

$$\begin{aligned}
&\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u,\theta_i)du \right. \\
&\qquad \left. - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}\right) \right) \mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}\right) \\
&+ \left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - (b_{i,k}-1)q_{i,k} \right) \mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}-1\right) \\
&\geq \left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k-1} \right) \mathbb{P}\left(P^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)=b_{i,k}\right) \\
&+ \left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r'_{i,k}} v^{g,h}(u,\theta_i)du \right. \\
&\qquad \left. - (b_{i,k}-1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r'_{i,k}\right) \right) \mathbb{P}\left(P^g\left(\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)=b_{i,k}-1\right) \qquad (13)
\end{aligned}$$

A.2 Upward deviation

Equilibrium of Model $S(0)$ also requires that the deviation of bidder $i \in \mathcal{N}^{g,h}, g \in \{A,B\}, h \in \{1,2\}$ at step k to a bid price of $b_{i,k}+1$ rather than $b_{i,k}$ weakly reduces her expected utility, and so requires

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)\right] \geq \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}''(\boldsymbol{\Theta})|\theta_i)\right] \\
&\implies \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\int_0^{y_i^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du - P^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))y_i^g(\boldsymbol{\sigma}(\boldsymbol{\Theta})) - c_i(K_i)\right] \\
&\qquad \geq \mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[\int_0^{y_i^g(\boldsymbol{\sigma}''(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du - P^g(\boldsymbol{\sigma}''(\boldsymbol{\Theta}))y_i^g(\boldsymbol{\sigma}''(\boldsymbol{\Theta})) - c_i(K_i)\right] \qquad (14)
\end{aligned}$$

By the Law of Iterated Expectations, $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}''(\boldsymbol{\Theta})|\theta_i)\right]$ can be decomposed analogously to $\mathbb{E}_{\boldsymbol{\Theta_{-i}}}\left[U_i(\boldsymbol{\sigma}(\boldsymbol{\Theta})|\theta_i)\right]$ in Equation 9.

First suppose that the bid prices are non-adjacent so that $b_{i,k}+1 < b^g_{i,k-1}$. I consider the case in which $b_{i,k}+1 = b^g_{i,k-1}$ in Appendix A.3.

By analogous reasoning to the case of the deviation to $(b_{i,k}-1)$, it follows that the probability of the event $P^g \in \{b_{i,k}, b_{i,k}+1\}$, and the allocation given each auction price $P^g \notin \{b_{i,k}, b_{i,k}+1\}$ are identical under $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta})$. Equilibrium therefore requires

$$
\begin{aligned}
&\left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - (b_{i,k}+1)q_{i,k-1}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}+1\right)\\
&+\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du\\
&\qquad - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)\right)\Bigg|P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}\Bigg]\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}\right)\\
&\geq\mathbb{E}_{\boldsymbol{\Theta}_{-\boldsymbol{i}}}\Bigg[\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g(\boldsymbol{\sigma}''(\boldsymbol{\Theta}))} v^{g,h}(u,\theta_i)du\\
&\qquad - (b_{i,k}+1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)\right)\Bigg|P^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)=b_{i,k}+1\Bigg]\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)=b_{i,k}+1\right)\\
&+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)=b_{i,k}\right) \qquad (15)
\end{aligned}
$$

Under Assumption 3.3 that the rationing coefficient is deterministic at each auction price, conditional on the strategy profile, equilibrium therefore requires

$$
\begin{aligned}
&\left(\int_0^{q_{i,k-1}} v^{g,h}(u,\theta_i)du - (b_{i,k}+1)q_{i,k-1}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}+1\right)\\
&+\Bigg(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u,\theta_i)du\\
&\qquad - b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}\right)\Bigg)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}(\boldsymbol{\Theta})\right)=b_{i,k}\right)\\
&\geq\Bigg(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r''_{i,k}} v^{g,h}(u,\theta_i)du\\
&\qquad - (b_{i,k}+1)\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r''_{i,k}\right)\Bigg)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)=b_{i,k}+1\right)\\
&+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du - b_{i,k}q_{i,k}\right)\mathbb{P}\left(P^g\left(\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right)=b_{i,k}\right) \qquad (16)
\end{aligned}
$$

A.3 Proposition 1 for adjacent bid prices

The following two conditions adapt Proposition 1 to the case in which bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ submits bids with adjacent bid prices on good $g \in \{A, B\}$.

As in Proposition 1, the equilibrium strategy profile is $\boldsymbol{\sigma}(\boldsymbol{\Theta})$, and $\boldsymbol{\sigma}'(\boldsymbol{\Theta})$ and $\boldsymbol{\sigma}''(\boldsymbol{\Theta})$ denote the strategy profiles which are identical to $\boldsymbol{\sigma}(\boldsymbol{\Theta})$ except that the bid prices submitted by type θ_i of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$ at step k, in her K_i-step bid function, are equal to $b_{i,k} - 1$ and $b_{i,k} + 1$, respectively.

Adjacent bid prices at steps $k-1$ and k only. In Model $S(0)$, under Assumption 3.3, in any type-symmetric Bayesian Nash Equilibrium, for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, every step $k \in \{1, ..., K_i\}$ in the K_i-step bid function in $\sigma_i(\theta_i)$ where $b_{i,k} - 1 = b_{i,k+1}$ and $b_{i,k} + 1 < b^g_{i,k-1}$, must satisfy

$$
\begin{aligned}
&\left(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u, \theta_i)du - b_{i,k} \left(q_{i,k-1} + (q_{i,k} - q_{i,k-1})r_{i,k}\right) \right) \mathbb{P}\left(P^g = b_{i,k} \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \\
&+ \left(\int_0^{q_{i,k}+(q_{i,k+1}-q_{i,k})r_{i,k}} v^{g,h}(u, \theta_i)du \right. \\
&\qquad \left. - (b_{i,k} - 1) \left(q_{i,k} + (q_{i,k+1} - q_{i,k})r_{i,k}\right) \right) \mathbb{P}\left(P^g = b_{i,k} - 1 \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \\
&\geq \left(\int_0^{q_{i,k-1}} v^{g,h}(u, \theta_i)du - b_{i,k} q_{i,k-1} \right) \mathbb{P}\left(P^g = b_{i,k} \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \\
&+ \left(\int_0^{q_{i,k-1}+(q_{i,k+1}} -q_{i,k-1})r'_{i,k} v^{g,h}(u, \theta_i)du \right. \\
&\qquad \left. - (b_{i,k} - 1) \left(q_{i,k-1} + (q_{i,k+1} - q_{i,k-1})r'_{i,k}\right) \right) \mathbb{P}\left(P^g = b_{i,k} - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \qquad (17)
\end{aligned}
$$

and Equation 5, where $q^g_{i,0} = 0$.

Adjacent bid prices at steps k and $k+1$ only. In Model $S(0)$, under Assumption 3.3, in any type-symmetric Bayesian Nash Equilibrium, for type $\theta_i \in \Theta_i$ of bidder $i \in \mathcal{N}^{g,h}, g \in \{A, B\}, h \in \{1, 2\}$, every step $k \in \{1, ..., K_i\}$ in the K_i-step bid function in $\sigma_i(\theta_i)$ where

$b_{i,k}-1>b_{i,k+1}$ and $b_{i,k}+1=b^g_{i,k-1}$, must satisfy

$$\begin{aligned}
&\Bigg(\int_0^{q^g_{i,k-2}+(q_{i,k-1}-q^g_{i,k-2})r_{i,k}} v^{g,h}(u,\theta_i)du \\
&\qquad\qquad -(b_{i,k}+1)\left(q^g_{i,k-2}+(q_{i,k-1}-q^g_{i,k-2})r_{i,k}\right)\Bigg)\mathbb{P}\left(P^g=b_{i,k}+1\big|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \\
&+\Bigg(\int_0^{q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}} v^{g,h}(u,\theta_i)du \\
&\qquad\qquad -b_{i,k}\left(q_{i,k-1}+(q_{i,k}-q_{i,k-1})r_{i,k}\right)\Bigg)\mathbb{P}\left(P^g=b_{i,k}\big|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \\
&\geq\Bigg(\int_0^{q^g_{i,k-2}+(q_{i,k}-q^g_{i,k-2}})r''_{i,k}v^{g,h}(u,\theta_i)du \\
&\qquad\qquad -(b_{i,k}+1)\left(q^g_{i,k-2}+(q_{i,k}-q^g_{i,k-2})r''_{i,k}\right)\Bigg)\mathbb{P}\left(P^g=b_{i,k}+1\big|\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right) \\
&+\left(\int_0^{q_{i,k}} v^{g,h}(u,\theta_i)du-b_{i,k}q_{i,k}\right)\mathbb{P}\left(P^g=b_{i,k}\big|\boldsymbol{\sigma}''(\boldsymbol{\Theta})\right) \qquad (18)
\end{aligned}$$

and Equation 4, where $q^g_{i,0}=q^g_{i,-1}=0$.

Adjacent bid prices at steps $k-1$, k and $k+1$. In Model $S(0)$, under Assumption 3.3, in any type-symmetric Bayesian Nash Equilibrium, for type $\theta_i\in\Theta_i$ of bidder $i\in\mathcal{N}^{g,h}, g\in\{A,B\}, h\in\{1,2\}$, every step $k\in\{1,...,K_i\}$ in the K_i-step bid function in $\sigma_i(\theta_i)$ where $b_{i,k}-1=b_{i,k+1}$ and $b_{i,k}+1=b^g_{i,k-1}$, must satisfy Equations 17 and 18.

B Proof of Corollary 1

At step $k=1$ of the K_i-step bid function in $\sigma_i(\theta_i)$ of type $\theta_i\in\Theta_i$ of bidder $i\in\mathcal{N}^{g,h}, g\in\{A,B\}, h\in\{1,2\}$, Equation 4 becomes

$$\begin{aligned}
&W\,\mathbb{P}\left(P^g=b^g_{i,1}\big|\sigma_i(\theta_i)\right)+X\,\mathbb{P}\left(P^g=b^g_{i,1}-1\big|\boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \\
&\qquad\geq Y\,\mathbb{P}\left(P^g=b^g_{i,1}\big|\sigma'_i\right)+Z\,\mathbb{P}\left(P^g=b^g_{i,1}-1\big|\boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \qquad (19)
\end{aligned}$$

where

$$W := \left(\int_0^{q_{i,1}^g r_{i,1}} v^{g,h}(u, \theta_i) du - b_{i,1}^g q_{i,1}^g r_{i,1} \right)$$
$$X := \left(\int_0^{q_{i,1}^g} v^{g,h}(u, \theta_i) du - (b_{i,1}^g - 1) q_{i,1}^g \right)$$
$$Y := 0$$
$$Z := \left(\int_0^{q_{i,1}^g r'_{i,1}} v^{g,h}(u, \theta_i) du - (b_{i,1}^g - 1) q_{i,1}^g r'_{i,1} \right)$$

Suppose $v^{g,h}(0, \theta_i) < b_{i,1}^g - 1$. Because the bidder's marginal valuation function is monotonic (Assumption 3.1.1), it follows that the bidder's marginal valuation function for good g is everywhere below $b_{i,1}^g - 1$. So, except in the knife-edge case in which $r_{i,1} = 0$, then $W < Y = 0$; and $X \leq Z \leq 0$.

As discussed in Appendix A,

$$\mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \geq \mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \geq 0 \tag{20}$$
$$0 \leq \mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) \leq \mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) \tag{21}$$

If $\mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) > 0$, then $W\, \mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) < Y\, \mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)$.

If $\mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) > 0$, then $X\, \mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) < Z\, \mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right)$.

So, providing $\mathbb{P}\left(P^g = b_{i,1}^g \middle| \boldsymbol{\sigma}(\boldsymbol{\Theta})\right) > 0$ or $\mathbb{P}\left(P^g = b_{i,1}^g - 1 \middle| \boldsymbol{\sigma}'(\boldsymbol{\Theta})\right) > 0$, and $r_{i,1} \neq 0$, these inequalities contradict Equation 19.

CHAPTER 3
Comparing the product-mix auction to alternatives

Abstract

This chapter compares the product-mix auction (PMA)—the mechanism used by the Bank of England (BoE) for its Indexed Long-Term Repo (ILTR) operations—to simpler alternative auction designs, namely a pair of separate simultaneous auctions, and a "reference price auction". Using data from the auctions held in June 2010 to January 2014, I find that the PMA increased welfare (defined by the difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept) by approximately 50%, or 2 basis points per loan, relative to these alternatives. I would expect larger welfare gains in a less stable period than the period studied, and simulations confirm this. Broader benefits of the auctions of reducing systemic risk, while mitigating moral hazard, informing the BoE about stress in the market, and communicating the "correct" price to the market, are taken into account in my approach, to the extent that the BoE's supply curve internalises some of these externalities. I also find that the PMA always gave the BoE more (or occasionally the same) surplus and revenue relative to if one of the alternative designs had been used. However, the effect of the PMA on aggregate bidder surplus was ambiguous. The latter result may be a property of the period studied, and of the fact that there were only two sets of eligible collateral in this period.

1 Introduction

The Bank of England (BoE) introduced the Indexed Long-Term Repo (ILTR) auctions in June 2010 to replace the long-term reverse repo operations with a mechanism that would efficiently respond to "the new demands for liquidity insurance that [the financial crisis] engendered" (Fisher, 2011a, p.15). The BoE used a novel design in which the quantity of funds lent against different types of collateral responded flexibly to the economic environment. The main objective of the design was to reduce systemic risk while avoiding moral hazard, and it also improved efficiency directly.

The ILTR auctions were designed to be conducted as product-mix auctions (PMA), which are single-round auctions designed to address the specific challenges of lending funds (i.e. central bank reserves) to financial institutions against multiple types of collateral. The maximum supply of funds is fixed, but the quantity of funds lent against collateral within a particular set varies depending on the bids received and the BoE's own supply preferences.[1,2]

This chapter assesses both the benefits of the design of the ILTR in terms of welfare gains, and the distribution of benefits between the BoE and bidders.

On the demand side, I assume that participants' bids express the prices they are willing to pay.[3] I measure the surplus of each participant by the difference between her bid and the price she pays (and sum this across bidders to get total bidder surplus). This does not take account of any externalities that result from meeting a participant's liquidity need and from charging it the correct price for doing so, but it equals her individual surplus.

On the supply side, my analysis relies on the mechanism being designed to maximise the overall benefits to society, that is, the BoE's own supply preferences express the "correct" combinations of prices and quantities. Any broader benefits (see below) of allocating funds

[1]Since 2014, the BoE does not commit to a maximum supply, and the total quantity of funds in the ILTR operations also responds to the bids and BoE's preferences. I therefore focus on the period June 2010 – January 2014 in which the same design was used for all auctions (except for some minor pricing changes, detailed in Section 2).

[2]Until October 2013, Level A and Level B collaterals were referred to as Narrow and Wide collaterals, respectively. They have also been respectively referred to as Strong and Weak collaterals. The Level A collateral set includes gilts, sterling Treasury bills and certain sovereign and central bank debt, whereas the Level B set includes less liquid sovereign debt and certain asset-backed securities. The BoE's current auctions also allow bids on (and generates a separate price for) loans of funds secured by collateral in a Level C set, which currently includes less liquid securities and portfolios of loans. For a full list of eligible collateral, see https://www.bankofengland.co.uk/markets/eligible-collateral.

[3]Chapter 1 finds that this appears to be a reasonable approximation in the ILTR auctions.

to bidders who are willing to pay a price above the BoE's supply curves are therefore taken into account to the extent that the supply curves internalise some of these externalities.

Importantly, I find that the welfare gain increases significantly when simulating a stress event, i.e. the benefits are more readily observed in times of higher liquidity demand.

To evaluate the PMA design, I compare alternative auction rules to the actual ILTR auctions that were held monthly in June 2010 – January 2014. The two main comparators are (i) separate simultaneous auctions (SSA) for loans secured by the different sets of collateral, each with a fixed supply of funds and (ii) a single-round "reference price auction" (RPA), in which the auctioneer fixes a pair of notional prices for the two collateral sets, i.e. "reference prices", and the highest bids relative to their reference prices are accepted. Setting reference prices in the RPA is equivalent to fixing the difference between the market clearing interest rates for the collateral sets in advance. I choose these two comparators as simpler reasonable alternatives, used either historically or by other central banks. For example, the BoE's gilt purchase operations, introduced in 2009, are run as RPAs, and the BoE's Corporate Bond Purchase Scheme, introduced in 2016, is a variant of an RPA. In my setting, I compare net welfare, bidder surplus, BoE revenue, and BoE surplus across auctions, for various fixed quantities (in the SSA), and fixed reference prices (hence, fixed price differences) (in the RPA). I also compare the outcomes of the PMA to separate sequential auctions and to posting fixed spreads over Bank Rate (with rationing if necessary) to borrow against the two collateral types.

The PMA jointly determines the quantities of funds lent against Level A and Level B collateral. It finds the competitive equilibrium, assuming that bids express bidders' true valuations for loans, i.e. bids are "truthful", and assuming that the BoE's supply curves express the minimum prices that it is willing to accept to supply loans secured by the collateral sets;[4] the difference between the market clearing interest rates ("prices") for funds lent against Level A and Level B collateral varies across auctions to implement these allocations. Under these assumptions, the PMA therefore maximises welfare (as conventionally measured as the sum of the BoE's and bidders' surpluses).[5] The alternative auction designs may fail to maximise welfare because of the constraints they introduce: the SSA and RPA do not permit quantities and the price difference, respectively, to adjust to the participants' preferences. Unlike

[4]Full details in Section 2.

[5]This is the case providing the clearing prices for the two goods are chosen to be the competitive equilibrium prices. The ILTR rules in the earlier part of the sample period permitted the clearing prices to occasionally differ from the competitive equilibrium, but this made little difference to the auction outcomes, and I describe the PMA as "welfare-maximising" throughout. See Appendix A for details.

the PMA, neither of these natural alternatives therefore allow a substantial increase in the allocation of funds against the less liquid Level B collateral in the event of stressed conditions, at prices the BoE is willing to accept. In Section 7, I consider alternative objectives to welfare, including maximising BoE revenue or surplus.

The main finding is that the flexibility in allocation of the PMA leads to welfare gains of approximately 50%, or 2 basis points per loan, as measured by the difference between the spreads that participants were willing to pay and the spreads that the BoE was willing to accept, relative to these two alternative designs. This is the case even though the period of analysis, June 2010 – January 2014, was a period without significant stress. Indeed, even with a hypothetical alternative objective of maximising BoE revenue or BoE surplus, rather than welfare, the PMA outperforms all possible RPAs and almost all possible SSAs. The benefits result from the automatic adjustment of the PMA allocation to the bids received and to the BoE's preferences between the two collateral sets. While the bidders were permitted to flexibly represent their preferences through "paired bids" for funds lent against Level A or Level B collateral over this period, this feature was rarely used.[6]

A second finding is that, in the period studied, the PMA always gave the BoE more (or occasionally the same) surplus and revenue relative to if the BoE had run any possible SSA or RPA, but the effect on the bidders, in aggregate, was ambiguous. It is trivial that the PMA must (weakly) increase aggregate surplus. However, it was slightly surprising that bidders, on average, shared in the gains only when the fixed share to Level B used in the SSA was low, or when the difference in fixed reference prices between A and B used in the RPA was high, so that the PMA typically lent more against Level B relative to the alternative. These results may be a property of the period studied, and of the fact that there were only two sets of eligible collateral, as I discuss below. I expect, and confirm in simulations, that the welfare gain and benefits to the BoE of the PMA would be qualitatively similar, but quantitatively larger in absolute terms, in a less stable period than the period studied. This reflects the scaling up of the benefits of the PMA's flexibility in the allocation and sensitivity to the BoE's preferences in contrast to the comparators. I also expect, as I explain and confirm in simulations below, that the specific features of the sets of eligible collateral mean that the PMA would benefit bidders over a larger range of SSAs, but a smaller range of RPAs, in a less stable period.

Importantly, none of the calculations take any explicit account of the broader benefits of

[6]Each bid of this kind specifies the quantity of loans demanded by the bidder and a bid price for each of the two collaterals. The bidder wins only if she would obtain positive surplus, and wins at most the quantity demanded of funds, borrowed against the collateral which gives her the highest surplus given her bids.

the PMA of reducing systemic risk while avoiding moral hazard (as well as informing the BoE about stress in the market and communicating the "correct" price to the market). The BoE wanted to permit banks to repo using collateral that was lower quality than the Level A collateral it had traditionally accepted, but only at prices which reflected the relative liquidity of the collateral assets. Moreover, the quantity of less liquid (Level B) collateral it was willing to accept was an increasing function of market stress. Clearly neither the SSA nor the RPA could achieve as fine-tuned a response as the PMA to market conditions. And so, because it does not capture all the broader benefits, my measure can be interpreted as a conservative estimate of the overall welfare gains of the PMA.

Section 2 provides background for the ILTR auctions and Section 3 describes the method. I compare the PMA to the SSA in Section 4 and discuss the impact of including reserve prices in the SSA, and of running the separate auctions sequentially. Section 5 compares the PMA to the RPA and to posting fixed prices. I expect the surplus comparisons would be more extreme in periods of stress, and Section 6 shows simulations that confirm this. Section 7 shows that alternative objectives to welfare-maximisation have only limited qualitative impact on the results. Section 8 describes the broader benefits of the PMA. Section 9 concludes.

1.1 Welfare results

SSA Had the BoE been able to forecast the optimal fixed quantities of funds to lend against Level A and Level B collateral separately, the SSA would have been almost as efficient as the PMA over the period studied. However, it seems unrealistic that the BoE could have identified this specific, best-performing SSA prior to the auctions' implementation. It introduced the PMA to overcome many operational policy questions relating to the existing LTR operations. In particular, the size of the operations and the appropriate spread between lending against Level A and Level B collateral were very unclear. The large increase in demand for funds in 2008 – 2009 added to this uncertainty. If the BoE had over- or under-estimated the optimal split of its total allocation of funds between the two separate auctions, the welfare loss incurred could have been substantial.

Importantly, the PMA allows the quantity allocated against Level B collateral to respond to stress. The composition of demand across Level A and Level B collateral in June 2010 – January 2014 turned out to be relatively stable. Had it not been, then any fixed split between the two collaterals in an SSA would have resulted in a significant welfare loss. For example, an unexpected increase in stress would have caused a potentially dramatic shift in

demand towards lending against Level B collateral.[7] Focusing on the period of June 2010 – January 2014 risks underestimating the advantages of the PMA relative to the SSA. When designing the auction, the BoE was seeking a mechanism that accommodated demand in periods of acute stress, rather than the relatively stable demand that is observed.

In particular, given that the objective of the ILTR was to lend funds against Level B collateral in stressed circumstances, any fixed split between Level A and Level B collaterals, chosen by the BoE, would presumably have allocated a significant proportion against Level B collateral. The results show that any fixed allocation against Level B exceeding 50% would have generated a welfare loss exceeding 2 basis points (bps) per loan of the quantity of funds in the ILTR operations and 35% of the welfare generated by the PMA.

RPA The difference in clearing prices between funds lent against Level A and Level B collaterals in the RPA equals the difference in their "reference prices", which are fixed in advance. There was enough variation in the optimal price difference between A and B across auctions that no single pair of reference prices came very close to maximising welfare in the period studied.

Across this period, the best-performing RPA would have generated an average welfare loss of 0.6bps and 0.8bps respectively per loan over a 3- and 6-month term, equivalent to 10% and 14% respectively of the average welfare generated by the PMA. Of course, this presumes the BoE would be able to choose the optimal price difference between collateral sets in advance. A reference price of B of 50bps and reference price of A normalised to zero (corresponding to the reserve prices set for Level B and Level A collateral in the auctions from October 2008 up until the June 2010 start of the PMA) would have almost maximised welfare for loans over a 6-month term, yet the average loss would have been around 3bps for loans over a 3-month term (equal to 53% of PMA welfare). In contrast, a reference price of B of 0bps (as in auctions prior to October 2008) would have generated an average loss for both the 3- and 6-month terms, of 2 and 5.3bps (32% and 71% of PMA welfare) respectively.

1.2 Distribution of surplus

The BoE always, and often significantly, benefited in both revenue and surplus (as measured by the difference between the market clearing prices for the loans and the prices the BoE

[7] This was observed in March 2020, in which the financial stress associated with COVID-19 led to a sharp increase in bids for funds lent against Level B collateral and especially against Level C, an additional, less liquid, collateral set introduced to the ILTR in 2014, after the sample period.

was willing to accept) from the PMA auction, relative to any possible SSA or RPA.[8,9]

However, perhaps surprisingly, bidders benefitted from the PMA relative to only some SSAs and RPAs.[10] Several features seem to have contributed to this.

First, the nature of demand for loans secured by Level A collateral is that it is always relatively flat at a price close to zero (and there is a reserve price of zero), so bidders can never make substantial gains or losses on Level A, whatever the auction form.

Second, crucially, in the time period in question, there was never very significant stress. As a result there was neither ever very large demand for loans secured by Level B, nor ever very large net surplus to bidders from borrowing against Level B (as bids were rarely much above the relevant point of the BoE's supply curve for Level B). Bidders would therefore have lost relatively little if the quantity lent against B had been below the welfare-maximising quantity, but would have gained substantially from very cheap borrowing against Level B if the quantity had been excessively large. Almost every SSA or RPA, fixed across the period, sometimes lent too little against B, and at other times too much, relative to the welfare-maximising PMA. The inefficiencies of the alternative designs benefitted the bidders in this respect.

A third feature was that bidders rarely expressed their preferences across the different collaterals, despite being able to do so via paired bidding.

Absent any one of these three features, the bidders might have benefited more from the use of the PMA.

If demand for loans secured by Level B collateral were more variable and sometimes much greater, the welfare benefits of the PMA relative to the SSA would likely have provided benefits to bidders, as well as the BoE. This is because the PMA would increase the allocation of the relatively valuable loans against Level B in response to the relatively high bids.

[8]There was one tiny exception over the period studied: running an RPA for 6-month loans with a fixed price difference of 50bps increases average BoE surplus by 0.01bps relative to the actual PMA.

[9]For any one auction, the PMA is, of course, identical to the SSA and RPA that are optimally calibrated for that auction, with the exception that the RPA may ration the bids incorrectly at the marginal prices. If demand is fairly stable, the PMA and an SSA that is optimally calibrated to the whole set of auctions are close to identical, but any given RPA may still lead to welfare losses, as discussed in Section 5.

[10]Chapter 4 shows that the PMA increases both bidder and auctioneer surplus relative to the best-performing SSA under reasonable conditions in a theoretical model in which bidders face correlated shocks to their values. It shows that these predictions are consistent with the estimated effect of using a PMA in Mexican Treasury auctions. These results suggest that both the bidders and the BoE would gain from the PMA. The smaller benefits I find for the bidders may be because shocks to bidders' values are partly idiosyncratic. Lings (2013) demonstrates that this reduces the likelihood that bidders gain from the PMA.

However, an RPA with a low price difference between Levels B and A would continue to largely prioritise bids on B, so that the PMA would not necessarily generate greater benefits for bidders. Some experimentation with different possible demands confirmed this claim (see Section 6).

If there was more than one collateral for which the demand curve was not flat, the PMA might provide more benefits to bidders. In particular, in a three-collateral implementation (Levels A, B and C, with C being the least liquid), such as the BoE currently uses, the bidders, as well as the BoE, might gain from the PMA, if the demand curves for two of the collaterals—presumably Levels B and C—are not flat.[11] This could be explored in further work.

And if bidders used the paired bids facility, they could directly gain from the use of the PMA (without loss to the BoE), in cases where they have the ability to substitute collaterals. Perhaps they did not have this ability, but it may be that they would in the future.[12] Alternatively bidders may not have been very sophisticated.[13] Again, their understanding could improve in the future.

2 Background

I study the ILTR auctions, conducted monthly by the Bank of England, over the period June 2010 – January 2014. Table 1 presents summary statistics.

The auction is a product-mix auction (PMA), described in Section 2 of Klemperer (2010, 2018).[14] In each auction, loans of central bank reserves for either a 3- or 6-month term are

[11] Bidders (as well as the BoE) would likely gain from the PMA relative to either using the SSA with a high share to Level C relative to Level B, or the RPA with a low price difference between Level C and Level B.

[12] As BoE Executive Director, Andrew Hauser, wrote (a little after the BoE decided it would stop using the paired-bids facility from February 2014, although its current programme is designed to accommodate paired bids), "higher collateral prices will give financial firms a strong incentive to invest in technologies allowing them to optimise their use of collateral. Indeed, big investment programmes are already underway in many firms and infrastructure providers, to ensure that firms have real-time information on the collateral they have available globally across all their business lines, that the collateral they deliver is cost effective, and that the cost of delivering (and financing) that collateral is factored into their risk and business decisions. These programmes involve sometimes relatively advanced technology; indeed, as some of our contacts remark, somewhat alarmed, 'for the first time in living memory, pointy heads are sitting on the repo desk'." https://www.bankofengland.co.uk/-/media/boe/files/speech/2013/the-future-of-repo-too-much-or-too-little

[13] Chapter 1 suggests that bidders were not too sophisticated in their use of the PMA, more generally — the good news is that this also means that strategic behaviour was unlikely (see Section 3.1).

[14] Klemperer (2010) is a minor revision of Klemperer (2008). Klemperer (2018) adds additional material to Klemperer (2010), but their Section 2s are identical.

Table 1: Summary statistics for ILTR auctions, June 2010 – January 2014

		3-month term		6-month term	
		Level A	Level B	Level A	Level B
Quantity of funds demanded (weighted, %)	Mean	79.69	20.31	66.44	33.56
	Std. Dev.	17.01	17.01	27.15	27.15
Quantity of funds allocated (weighted, %)	Mean	83.28	16.72	67.46	32.54
	Std. Dev.	13.41	13.41	24.56	24.56
Price paid (weighted average, bps)	Mean	1.17	22.12	0.9	38.91
	Std. Dev.	1.74	6.8	0.86	16.7
Bid price (weighted, bps)	Mean	2.5	22.98	3.06	35.87
	Std. Dev.	2.85	10.28	4.31	17.14

auctioned at a spread over Bank Rate.[15]

Within each auction, bidders may borrow the reserves against one of two types of collateral: Level A, including gilts, sterling Treasury bills and certain sovereign and central bank debt, and Level B, including less liquid sovereign debt and certain asset-backed securities.

Bids Bidders can submit any number of sealed bids, each of which can be either of two types:

1. Specific bids: Each bid of this kind specifies the type of collateral the bidder will provide, the quantity of funds that she demands, and the spread that she is willing to pay for the funds borrowed against the specified type of collateral. For example, a bid may specify demand of £50 million for a spread of 2bps against Level B.

2. Paired bids: Each bid specifies the quantity of funds the bidder demands, and the spread that she is willing to pay for the funds borrowed against each of the Level A and B collaterals. For each bid, she may be allocated the demanded funds against *either* Level A *or* Level B collateral.[16] For example, a bid may specify demand of £50 million for a spread of 2bps against A or 25bps against B. The bid will be allocated up to £50 million against *either* Level A *or* Level B.

Paired bids In practice, bidders rarely submit paired bids so I drop these from the sample. This change is insubstantial.

[15]Loans in the current ILTR auctions are for 6-month terms. The Short-Term Repo (STR) facility (announced in August 2022) will run alongside the ILTR auctions and offer loans with one-week maturity against Level A collateral at a fixed at Bank Rate (with no maximum supply).

[16]If both paired bid prices are higher than the respective clearing prices, the bid is allocated to the good that generates the highest surplus to the bidder (determined by the difference between the bid and clearing price).

The minimum bid size is £5 million, with increments of £1 million, and the minimum unit of allocation is £100,000. There is no obligation for registered bidders to participate in an auction, and bidders can bid up to 30% of the maximum supply (£1.5 billion and £750 million in the 3- and 6-month term auctions respectively).

Supply Prior to any one auction, the BoE commits to two supply curves, which are illustrated in Figure 1:

1. A maximum supply: The BoE announces a maximum supply made available across both Level A and B collaterals. The maximum supplies were £5 billion in each of the 3-month term auctions and £2.5 billion in each of the 6-month term auctions.[17]

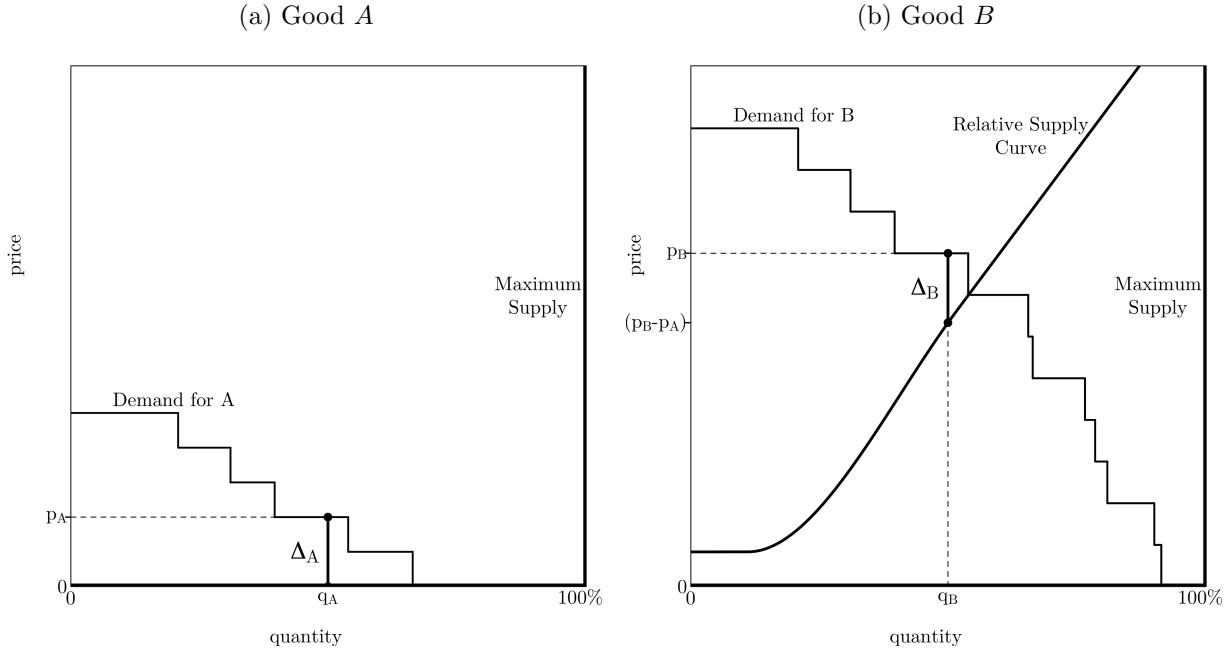
2. A "relative supply" curve: The BoE commits to a privately known "relative supply" curve. This represents the minimum price that the BoE is willing to accept on the margin to supply funds against less liquid Level B collateral relative to Level A collateral.[18] The BoE's supply preferences are therefore measured in terms of the difference between the spread for Level B collateral and the spread for Level A collateral, as a function of the percentage of the maximum supply allocated against Level B collateral. This influences the welfare-maximising allocation across the two collateral sets.

I refer to funds borrowed against each type of collateral as goods A and B respectively. The market clearing spreads over Bank Rate for loans of funds secured by each set of collateral are referred to as the clearing prices for goods A and B.

[17]Since 2014, the BoE has committed to an upward-sloping "total supply", which is a function of the total quantity allocated across all collateral sets, rather than a maximum supply. The maximum supply in my context is equivalent to a total supply curve which is flat at zero up to the maximum supply, at which point it becomes vertical.

[18]The difference in the prices that the BoE is willing to accept to lend against Level B and A collaterals reflects differences in liquidity. (The BoE imposes a haircut to each collateral asset to account as far as possible for differences in risk, see de Roure and McLaren (2021).)

Figure 1: Competitive equilibrium prices, (p_A, p_B), and allocations, (q_A, q_B), in the PMA. The welfare gain from allocating a marginal unit is equal across the two goods, i.e. $\Delta_A = \Delta_B$



Market clearing rules The PMA is a uniform-price auction, so that all winning bids pay the single market clearing price for the good they win. As described in Klemperer (2008), the original proposal was that these two prices (one each for goods A and B) were to be chosen to be competitive equilibrium prices. This would mean that the auction would maximise social welfare (the sum of BoE and bidder surplus), assuming bidders bid their actual valuations, which is a reasonable assumption.[19] To implement this, the PMA finds the quantities such that the difference between the clearing prices for goods B and A is equal to the difference in the prices the BoE is willing to accept on the margin between B to A. Figure 1 shows the competitive equilibrium prices, and the corresponding quantities allocated, for an illustrative realisation of bids. It is easy to see from the figure that either increasing or reducing q_A (and changing q_B correspondingly) would reduce total welfare.[20] See Section 3.1 and Appendix A for discussion.

However, even where the competitive equilibrium prices are unique (they generally are—see Appendix A.1) there is often a range of prices at which demand (although not the BoE's

[19] At least for bidders who are not bidding for too large a fraction of the amount offered, bidding truthfully is (close to) optimal.

[20] Total welfare is measured as the sum of the area under the aggregate demand curve for good A up to q_A and the area between the aggregate demand curve for good B and the relative supply curve up to q_B (see Section 3.3).

Table 2: PMA market clearing rules for the ILTR auctions, June 2010 – January 2014

Operation Date	Price Determination		Relative Supply Curve	
	Good (Level) A	Good (Level) B	3-month term	6-month term
Jun 2010 – Apr 2011	Lowest-winner pricing	Lowest-winner pricing	Starts at 5bps	Starts at 5bps
May 2011 – Aug 2011	Lowest-winner pricing	Lowest-winner pricing	Starts at 5bps	Starts at 15bps
Sep 2011 – Apr 2012	Lowest-winner pricing	Highest-loser pricing	Starts at 5bps	Starts at 15bps
May 2012 – Jan 2014	Highest-loser pricing	Highest-loser pricing	Starts at 5bps	Starts at 15bps

offered supply) matches the competitive equilibrium quantities. Appendix A describes the precise price determination rules, and Table 2 summarises them as they vary over the period. On both A and B, the rules changed from "lowest-winner pricing" (LWP) to "highest-loser pricing" (HLP).[21] The final rules—HLP on both A and B—correspond to the rules originally proposed in Klemperer (2008).[22] I call these rules PMA-HLP.

Bids strictly above the clearing prices are fully allocated, and bids on the margin for a good are rationed pro-rata.

Table 2 also shows the starting value of the relative supply curve as it varies over the period, separately for the 3-month and 6-month term auctions. The relative supply curve is an increasing step function, with step length equal to 0.1% of the maximum supply.

In Table 2, the rule changes are shaded in grey.[23]

[21]Under LWP, the price for each good is equal to the lowest winning bid on that good. Under HLP, the price for good A is equal to the highest losing bid on that good, or 0bps if all bids on good A are fully allocated. The price for good B under HLP equals the clearing price for good A plus the value of the relative supply curve at the quantity allocated of good B.

Using LWP introduced slight welfare losses in the earlier period under specific circumstances. When the auction was not fully allocated but the clearing price of good A was strictly positive, the implied minimum price that the BoE was willing to accept on the margin for good B was raised above the true price that the BoE was willing to accept. In four auctions, this caused the quantity allocated of good B to be slightly below the optimal quantity, so welfare was not maximal.

To evaluate the significance of this change, I repeat the surplus difference calculations using HLP for the PMA price determination in all auctions. The results are very similar to the corresponding results for the actual PMA, shown in Appendix A.4. The net welfare differences are almost identical and the transfers from bidders to the BoE under the PMA are slightly smaller, as the clearing prices are on average lower when using HLP.

[22]HLP is preferable to LWP as small bidders, at least, then have no incentive to misrepresent their preferences. If a bidder expects LWP, she has an incentive to underreport her true demand, due to the possibility of her bid determining the market clearing price. Regardless of the size of her demand, a bidder's incentive to distort her bid away from her valuation is larger under LWP than HLP.

[23]In many cases, the changes in price determination rules are irrelevant. However, in cases in which the auction is uncovered, the highest losing bid on good A is interpreted as equalling 0 (i.e. the reserve price), which may be lower than the lowest winning bid. Similarly, the price that the BoE was willing to accept on the margin for good B may be lower than the lowest winning bid.

Usage Usage of the auctions substantially decreased from June 2012.[24] The greater demand for funds in the earlier auctions increases the scope for misallocation and welfare losses if the operation is not appropriately designed. The average differences in surpluses between designs are weighted by the quantity allocated in the PMA.

3 Method

3.1 Truthful bidding

I assume that bidders submit bids that describe their true willingness to pay for the two goods. If bidders believe themselves to be price takers, so that their bids do not affect the market clearing prices, bidding truthfully is an optimal strategy.

This assumption is motivated by the observation that a large proportion of bidders submit a single bid, constituting a flat demand curve. If bidders expect to be pivotal and act strategically, bidders optimally submit multiple staggered bids (under reasonable assumptions on bidders' preferences), and therefore engage in "laddering".[25] Further, the number of potential participants is known to be large,[26] and neither the actual number of participants nor the BoE's relative supply curve for any auction is ever disclosed to the bidders,[27] making optimal strategic behaviour difficult to determine.[28] Chapter 1 finds that truthful bidding

[24]The average cover in June 2012 – January 2014 is 0.4% for good A and 3% for good B, compared to 74% and 23% respectively in June 2010 – May 2012. A plausible reason for this is an increase in the availability of funding from other sources. The Funding for Lending Scheme was introduced by the BoE in July 2012, allowing borrowing of up to four years. In addition, the BoE's Extended Collateral Term Repo (ECTR) operation, consisting of monthly auctions for 6-month loans is another source, although this was only activated once (in June 2012). Increased availability of funds from the euro area also plausibly reduced demand in the ILTR auctions.

[25]A bidder with market power has an incentive to misreport her preferences. By understating her willingness to pay, she pushes down the expected market clearing price and reduces her expected total payment for the allocation of funds. A fully strategic bidder trades off this gain with the possibility of also losing in cases in which she would prefer to win.

[26]Participants of the BoE's Sterling Monetary Framework with access to the BoE's Open Market Operations (OMO) are eligible to register to participate in the ILTR auctions. The number of institutions with OMO access (i.e. the number of potential participants) was fairly stable from October 2009 (before the ILTRs were introduced) to January 2014 and ranged from 48 to 52.

[27]Bidders could have inferred some features of the supply curve from the information publicly released after the auctions. For example, in February 2011, the 6-month term auction is uncovered but some bids on good B are unallocated and the marginal bid on good B is rationed. Given this, the bidders could infer a point on the 6-month term relative supply curve: at the total quantity allocated of good B (33.3%), the BoE's relative supply equals the difference in the clearing prices for B and A (52bps). However, the amount of information revealed over time was limited and the supply curves shifted over time as shown in Table 1, limiting the scope to act on this information strategically.

Bidders could also have inferred the number of participants to be small in some of the later auctions (from October 2012 onwards), based on the publicly revealed information.

[28]Because bidders for both goods bid within the same auction in the PMA, the PMA increases competition

appears to be a reasonable approximation in the ILTR auctions.

Under this assumption of truthful bidding, the PMA maximises welfare.

3.2 Alternative auction rules

I use the PMA as a welfare-maximising benchmark to compare two alternative auction rules.

3.2.1 Comparison 1: Separate simultaneous auctions

Separate simultaneous auctions with fixed quantities supplied (SSA) I first compare the PMA to running two separate auctions simultaneously, one for each good, with the quantity supplied of each good, A and B, announced and fixed prior to the auction.

For each of the 3-month and 6-month term auctions, I fix the total supply available across both goods to the maximum supplies of £5 billion and £2.5 billion, respectively. I then vary the fixed percentage of this total that is supplied of good B from 0% to 100%. Under each of these auction rules, I calculate the market equilibrium with reserve prices equal to the starting point of the supply curve in each period, shown in Table 1, assuming that bidders' bidding behaviour remains unchanged. The market clearing prices equal the highest losing bids for the two goods. I then compare the realised surpluses and net welfare relative to the welfare-maximising PMA benchmark.

The SSA constrains the BoE relative to the PMA, eliminating the flexibility in supply that the PMA allows both through the joint determination of clearing prices and through the relative supply curve.

Importantly, the SSA allocation does not depend on the relative supply curve, and so there may be outcomes in which the BoE is supplying good B at a price below what it would be willing to accept on the margin. Large welfare losses can therefore arise in SSAs in which the fixed quantity allocated of good B is large, for auctions with high demand for good B.

Reserve prices To test for robustness of the results, I also consider variants of the SSA, in which there are (i) no reserve prices, (ii) reserve prices of 5bps for good B across the whole period and (iii) reserve prices of 5bps for all 3-month auctions and 15bps for all 6-month auctions. The SSA that I focus on is the most favourable one in terms of both welfare and BoE surplus. In this design, the reserve prices change over time in line with the observed

relative to the SSA, which may reduce the incentive to bid strategically.

changes in the actual supply curve, despite the BoE making these changes in response to signals from the PMA that would have been unobserved in an SSA. The welfare losses from this SSA can therefore be thought of as a lower bound on the welfare losses resulting from an SSA.

I also discuss the impact of running the two separate auctions sequentially.

3.2.2 Comparison 2: Reference price auction

Reference price auction (RPA) Second, I compare the PMA to running a single-round auction in which a pair of notional prices for goods A and B are fixed in advance, and the highest bids across the goods relative to their reference prices are accepted, with the total supply available across both goods fixed and announced prior to the auction. This design has been used in financial markets, for example, the BoE and Federal Reserve have used RPAs for their quantitative easing purchase schemes (Song and Zhu, 2018; Bank of England, 2022).[29] The notional prices might reasonably be fixed at "best guesses" of the fair market prices. For example, the BoE's gilt purchase operations and Corporate Bond Purchase Scheme are RPAs in which the reference price for each bond is its secondary market mid yield at the close of the auction (Bank of England, 2022).

To implement an RPA, I normalise the reference price of good A to zero, subtract the reference price of B from the bids on good B, and find the market clearing allocation as if all bids are for a single good. This then implies clearing prices for goods A and B, equal to the highest losing bids for the two goods.[30] The fixed reference prices mean that the difference in clearing prices for goods A and B equal the difference in reference prices. The marginal winning bid is restricted to be non-negative. This effectively imposes a reserve price for good A of 0bps and a reserve price for good B equal to its reference price.

Given the nature of the two "goods"—loans of funds secured by Levels A and B collateral—this can be interpreted as imposing a fixed spread between loans secured by collateral assets with different liquidity levels.

I also compare the PMA to posting a pair of fixed prices, with the total supply available

[29] Armantier, Holt and Plott (2013) report that the US Treasury also considered using an RPA to purchase "toxic assets" (namely mortgage-backed securities) in the $700 billion Troubled Asset Relief Program (TARP) in 2008.

[30] Bids strictly above the respective clearing prices are fully allocated and bids on the margin in a fully allocated auction are rationed pro-rata with a common rationing coefficient across goods A and B.

fixed and bidders with bids above the posted prices rationed pro-rata. This can be viewed as an RPA in which the auction prices are also fixed.

3.3 Surpluses, and welfare

I calculate the auction outcomes under the different rules and calculate the difference in total net bidder surplus (denoted bidder surplus), BoE surplus, BoE revenue, and net welfare between the actual PMA and the alternative auction.

Bidder surplus A bidder's bids for a particular good can be aggregated to form her individual demand curve for that good, which is equal to her marginal valuation curve under the assumption of truthful bidding. For each bidder, net bidder surplus therefore equals the sums of the areas under her demand curves for goods A and B up to the quantities allocated, minus the total payment that they make. Bidder surplus, equal to total net bidder surplus, is the sum of individual net bidder surpluses.

BoE surplus The maximum supply in each auction and the position of the relative supply curve are "pinned down by [the BoE's] preferences — namely to provide liquidity insurance to the banking system at a price that doesn't undermine the incentive to prudently manage liquidity" (Fisher, 2011a, p.12). I interpret this to mean that the supply curves reflect the minimum prices that the BoE is willing to accept for the loans, as functions of the allocation. In particular, for any total quantity less than the maximum supply, the price that the BoE is willing to accept for good A is zero and the price that it is willing to accept to supply a marginal unit of good B equals the height of the relative supply curve. I also assume that, for each auction, the price at which the BoE is willing to supply any quantity in excess of the maximum supply is infinite, so the total quantity allocated can never exceed the maximum supply.[31] For good A, BoE surplus is therefore the market clearing price for good A multiplied by the quantity allocated of good A. For good B, it is the difference between the market clearing price for good B and the BoE's relative supply curve, up to the quantity allocated of good B. See Appendix A.3 for further discussion.

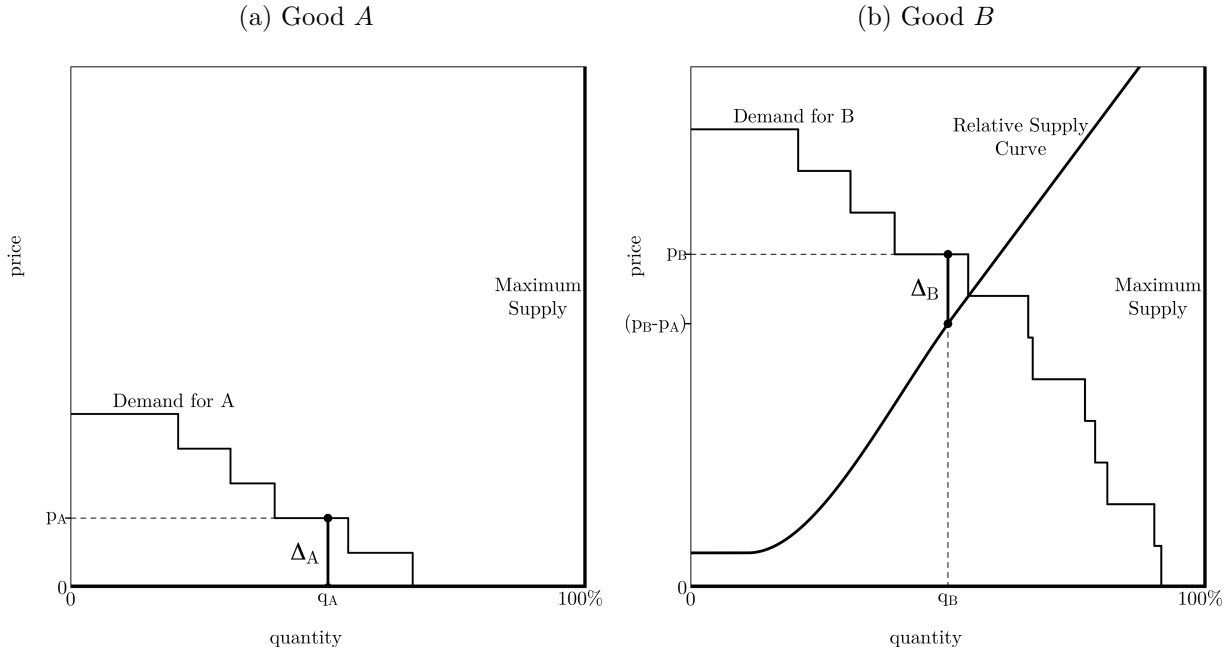
BoE revenue BoE revenue is measured by the sum of the quantity allocated of good B multiplied by the market clearing price for good B plus the quantity allocated of good A multiplied by the market clearing price for good A.

[31] The relative supply function is a step function, over a fine quantity grid, with steps at each 0.1 percentage point of the maximum supply. For quantities allocated between the two steps, the price that the BoE is willing to accept equals the lower bound of the two steps.

Net welfare Total net welfare equals the sum of bidder surplus and BoE surplus.[32,33]

Figure 2 illustrates these surpluses for each good in the PMA, at market clearing prices (p_A, p_B) and quantities (q_A, q_B). Surpluses in the alternative designs are analogous.

Figure 2: Surpluses at the PMA market clearing prices, (p_A, p_B), and quantities, (q_A, q_B)



4 The SSA

I first compare the surpluses resulting from the actual PMA and the SSA, as the proportion of total quantity supplied of good B varies from 0% to 100%.

Figure 3 show the differences in auction surpluses between the actual PMA and SSA, averaged over the 3-month and 6-month term auctions in June 2010 – January 2014, weighted

[32]The welfare-maximising total quantity allocated may be less than the maximum supply even if the auction is "covered" (i.e. total demand across both goods at least equals the maximum supply). The reserve price for good A is 0bps, and the relative supply curve constrains the welfare-maximising allocation of good B so that the difference between the lowest winning bids for goods B and A weakly exceeds the difference in prices that the BoE is willing to accept on the margin for goods B and A (equal to the height of the relative supply curve).

[33]To illustrate, consider the July 2010 auction. In this auction, the clearing prices for goods A and B are 1bp and 26bps, respectively. 17.2% of the maximum supply is allocated of good B, and 82.8% allocated of good A. The marginal bid on good A is 1bp, yielding marginal net welfare of 1bp. The marginal bid on good B is 26bps and the price that the BoE is willing to accept on the margin for good B (equal to the relative supply curve evaluated at 17.2%) is 25bps. So the marginal net benefit of allocating a unit to good B is also 1bp. Equating the two marginal net benefits maximises net welfare.

by the quantities allocated in the PMA. The surpluses in an auction are measured in basis points of the quantity allocated by the PMA in that auction.

There are five key features of the results.

Feature 1 Welfare maximisation could almost have been achieved in the SSA given observed demand

In Figure 3, the average net welfare difference between the SSA and PMA is U-shaped, strictly decreasing to approximately zero and then strictly increasing. This shows the welfare loss of imposing the SSA constraint on the auction. If the quantities supplied of good B were fixed at 17.5% for 3-month terms and 25% for 6-month terms, the average welfare loss from running the simpler SSA would have been essentially zero.

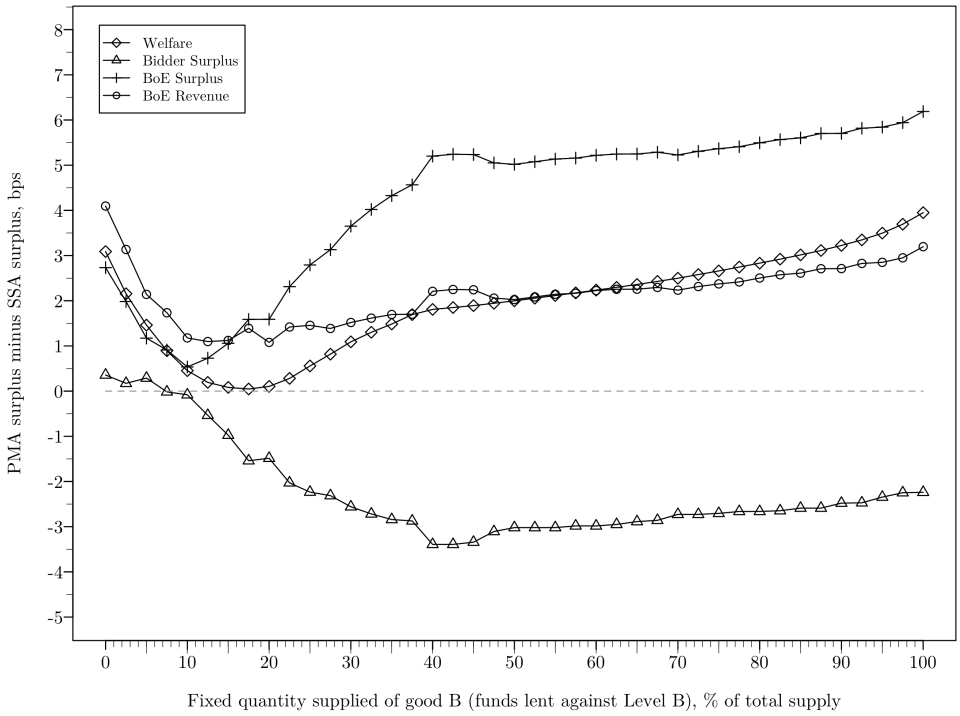
This suggests that, had the BoE been able to forecast these optimal quantities, implementing the SSA would have almost maximised welfare over the period. However, this seeming strength of the SSA is limited for two reasons:

1. **Forecasting** These optimal fixed quantities were very uncertain ex ante (Fisher, 2011a, p.8), so the BoE could not have implemented the best-performing SSAs in practice and the welfare losses of forecast error can be significant (see Feature 2). Moreover, if the BoE had adjusted the fixed quantities based on the outcomes of the initial auctions, the welfare losses might have been small over the sample period because the optimal fixed quantities turned out to be fairly constant. However, it could not have relied on this stability ex ante. (This uncertainty in both the optimal quantities and their stability was one motivation for the introduction of the ILTR auctions, replacing the LTR operations, in 2010.)

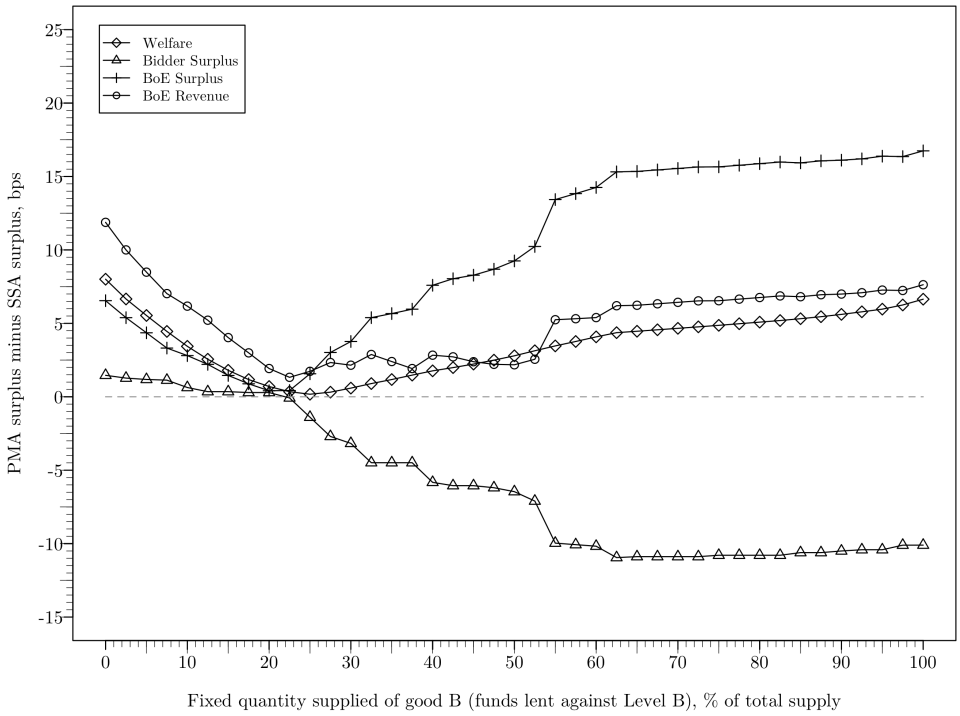
2. **Endogenous response** The ILTR auctions were introduced and designed to flexibly and systematically respond to financial stress. The composition of demand for funds in the June 2010 – January 2014 period was relatively stable, so that the optimal quantity supplied of good B was relatively stable across auctions. In a time of stress, demand for good B is likely to rise relative to good A, so that the optimal quantity allocated of B would likely rise. If unanticipated, an SSA could generate significant welfare losses relative to the PMA, and if anticipated, a publicly announced change in the composition of the SSA could add to market instability.

Figure 3: Difference in auction surpluses, PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term



Feature 2 But the welfare losses could have been large if the BoE had misjudged the correct proportion to assign to good B

Figure 3 shows that the average net welfare loss from the SSA can be significant. For example, if the BoE had fixed the quantity supplied of good B to be 50% of the maximum supply, this would have led to an average welfare loss of over 2bps per auction for both 3- and 6-month terms, equal to 32% of the average welfare generated by the PMA: if the BoE had fixed the quantity supplied of good B to be 2/3 of the maximum supply, the average welfare losses would have been approximately 2.4bps and 4.6bps (41% and 47% of PMA welfare) respectively for the 3- and 6-month term auctions.[34]

Feature 3 Net welfare changes smoothly but changes in bidder surplus and BoE surplus and revenue are lumpy

The average difference in net welfare changes smoothly as the fixed quantity supplied of good B rises. In contrast, the average revenue, bidder surplus and BoE surplus differences are lumpy.

The allocation changes approximately continuously[35] as I change the proportion allocated to good B, but the clearing prices are discrete and so change discontinuously. In addition, at the optimum of 17.5% for 3-month terms and 25% for 6-month terms, welfare losses are second order whereas the transfers between agents are first order. Net welfare is determined by the allocation, whereas bidder surplus, BoE surplus and revenue are determined by both the allocation and clearing prices.

Feature 4 Large transfer from BoE to bidders in some SSAs relative to the PMA

In Figure 3, the average BoE surplus is uniformly higher in the PMA than the SSA across all fixed quantities supplied of good B. The average bidder surplus is higher in the PMA than the SSA for low fixed quantities supplied of good B (up to 7.5% fixed quantity in the

[34]Observe also that if none of the maximum supply were allocated of good B, the average welfare losses would have been 3.1bps and 8bps (54% and 66% of PMA welfare) respectively for the 3- and 6-month term auctions. This gives some indication of the foregone welfare had the BoE not extended the acceptable range of collateral to include Level B collateral in the liquidity operations from June 2010. (This estimate is an upper bound), because some bidders who bid against Level B collateral in the PMA might perhaps have submitted bids against Level A, had Level B collateral not been accepted. However, the rare use of paired bids suggests this substitution behaviour might have been limited.)

[35]The smallest unit of allocation is £100,000, equal to 0.002% and 0.004% of the maximum supply in the 3-month and 6-month term auctions respectively, so the allocation can be treated as approximately continuous.

3-month auctions and 20% in the 6-month auctions). It is uniformly higher in the SSA than the PMA across all fixed quantities higher than these respective thresholds. In particular, average bidder surplus is 1.5bps higher in the best-performing SSA than the PMA for the 3-month auctions and marginally higher for the 6-month auctions. This is the case despite the fact that the efficiency loss of the best-performing SSA is small. The allocation chosen by the PMA is not constant, and small differences in the allocation between the PMA and best-performing SSA can create large changes in bidder and BoE surpluses through changes in the clearing prices.

There are two reasons for these differences in bidder surplus and BoE surplus:

1. Misallocation of goods A and B
2. Transfer of surplus between the BoE and bidders through differences in the market clearing prices

First consider the SSAs with very low fixed quantities supplied of good B and corresponding high fixed quantities supplied of good A. In these cases, good A is over-allocated and good B is under-allocated relative to the optimal quantities. There is large misallocation in the SSA, causing both bidder surplus and BoE surplus to be higher in the PMA. The high fixed quantity allocated of good A also pushes down the clearing price for good A, yielding a transfer from the BoE to bidders in the SSA.

For higher fixed quantities supplied of good B, bidder surplus is lower and BoE surplus is higher in the PMA relative to the SSA. As the fixed quantity allocated of good B rises beyond a certain point, good B is over-allocated and A is under-allocated, holding the total quantity fixed. In addition, the SSA allocation is determined independently of the BoE's relative supply curve so that there may be units allocated of good B for which the price that the BoE would be willing to accept exceeds the marginal benefit of allocating the loans. This generates further loss in BoE surplus and gain in bidder surplus in the SSA relative to the PMA.

The transfer of surplus between the BoE and bidders in the SSA may also be substantial at these higher fixed quantities supplied of good B. The market clearing prices in the SSA are determined entirely by the highest-losing bids, independent of the prices that the BoE would be willing to accept for the marginal loan. As the fixed quantity supplied of good B rises, eventually all of the bids on good B are fulfilled, so that the clearing price for good B falls to the reserve price. In contrast, the market clearing price in the actual PMA either

equals the lowest winning bid or the maximum of the highest losing bid and the price the BoE is willing to accept for the marginal loan. The market clearing price for good B in the SSA is typically far below that in the actual PMA, generating a large transfer from the BoE to the bidders.

In addition, the transfer from BoE to bidders is partly driven by the price determination rules. Recall the PMA moves from an LWP rule to an HLP rule (which is a function of the price that the BoE is willing to accept on the margin). In contrast, the SSA adopts an HLP rule throughout the entire period, which is independent of the supply curve. By construction, the LWP increases BoE surplus for a given allocation (assuming truthful bidding). To investigate this, I also repeat the welfare calculations for a PMA with HLP for the entire period, where the highest loser in the PMA case is the maximum of the price that the BoE is willing to accept on the margin and the highest losing bid. The differences in BoE surplus and bidder surplus are slightly attenuated, especially around the range of minimum welfare loss.[36]

Feature 5 Reserve prices do little to reduce misallocation

One key source of welfare loss in the SSA is an allocation of good B in excess of the optimal quantity so that the price that the BoE would be willing to accept on the margin exceeds the marginal bid (equal to the marginal bidder valuation).

A crude way to mitigate this welfare loss is to set reserve prices in the SSA. I assumed the SSA has reserve prices equal to the minimum spreads of the relative supply curve in each auction, shown in Table 1. I also calculate the welfare losses under three different reserve prices: (i) no reserve prices, (ii) reserve prices of 5bps for good B across the whole period, and (iii) reserve prices for good B of 5bps for all 3-month auctions and 15bps for all 6-month auctions. The average surpluses for 3- and 6-month auctions are very similar in all four cases.

4.1 Separate auctions run sequentially

Running separate auctions sequentially would almost certainly yield similar results to running them simultaneously, because the fact that paired bids were hardly used suggests that bidders would bid the same way in each auction whether they were sequential or simultane-

[36]The benefit to bidders of the PMA is small (but positive) when comparing the PMA-HLP to an SSA with a small quantity supplied against Level B, but the bidders receive essentially no benefit from the PMA when comparing with the actual version of the PMA used (with LWP in the earlier auctions, rather than HLP).

ous.[37,38]

5 The RPA

Second, I compare the surpluses between the PMA and the reference price auction (RPA) as the reference price of B varies from 0bps to 60bps. The RPA seems a particularly natural comparator since the BoE uses the design for its quantitative easing purchases and used a somewhat related design in the LTR auctions prior to the implementation of the ILTR auctions in June 2010.[39]

Figure 4 show the differences in auction surpluses between the PMA and RPA, averaged over the 3-month and 6-month term auctions in June 2010 – January 2014, weighted by the quantities allocated in the PMA. As for the SSA comparison, the surpluses in an auction are measured in basis points per loan of the quantity allocated by the PMA in that auction.

There are five key features of the results.

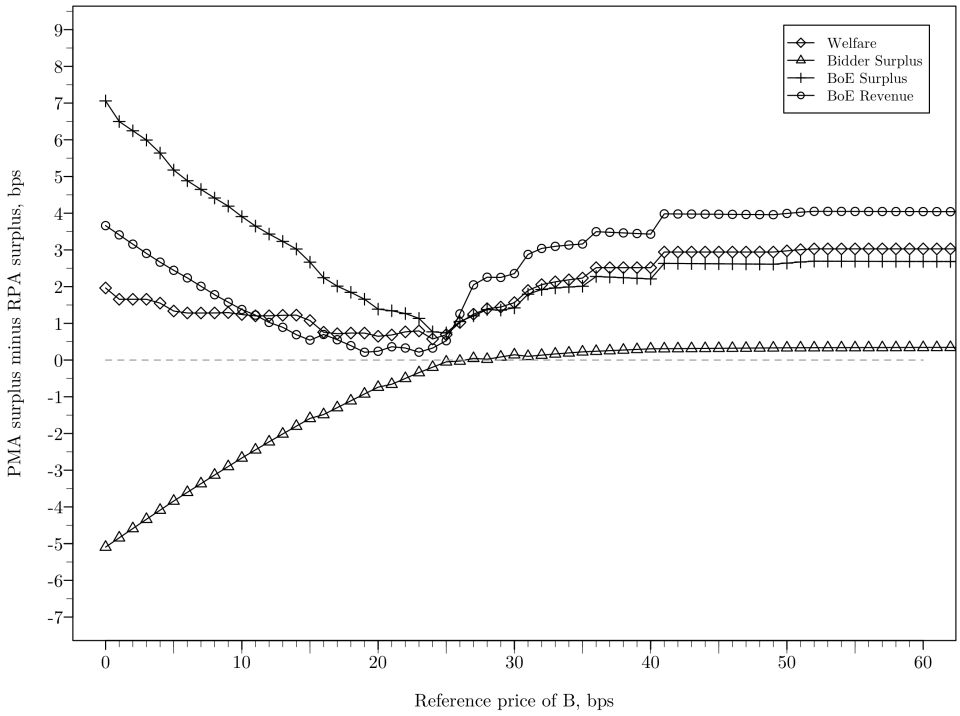
[37]Paired bids allowed bidders to express one-for-one substitute preferences, but bidders with more complex substitute preferences could still have gained by using them, so the near-absence of paired bids suggest that bidders viewed the goods independently. Moreover, either one-for-one substitutability or non-substitutability between the goods appears most plausible: they are measured in the same units and generate the same amounts of liquidity. Indeed, the BoE designed the ILTR auction anticipating bidders to express one-for-one substitutability, if at all.

[38]Running the auctions sequentially permits some adjustment between auctions by either bidders or the auctioneer. The information bidders learn in the first auction could affect behaviour in the second auction, but the effect seems likely to be small. The auctioneer could somewhat improve the performance of the sequential auctions by making any unused supply from the first auction available in the second auction.

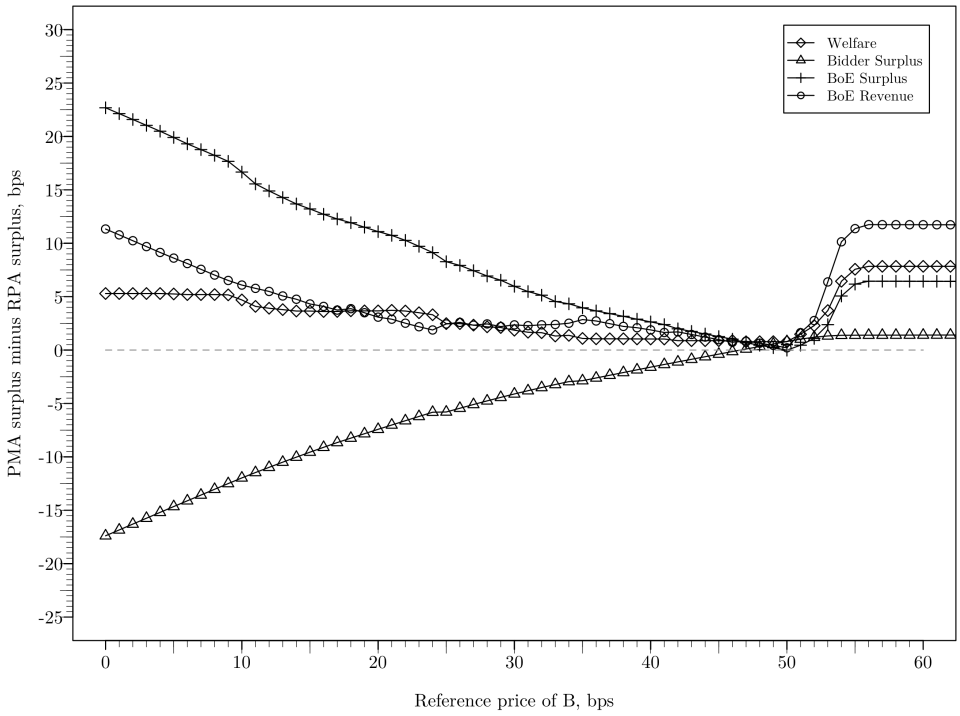
[39]In the LTR operations, loans over a 3- to 12-month term were auctioned against both Level A and B collateral sets. This was a discriminatory auction in which bids were accepted in decreasing order, with a reserve price of 50bps applied exclusively to bids for loans secured by Level B collateral.

Figure 4: Difference in auction surpluses, PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term



Feature 1 Welfare could not have been maximised with RPAs given observed quantities

The average net welfare difference between the actual PMA and the RPA does not fall to zero in either the 3-month or 6-month auctions over the period studied. Figure 4 shows that, for the entire range of reference prices of B, the average welfare loss of the RPA is at least 0.6bps and 0.8bps, respectively for the 3- and 6-month auctions. This is equivalent to 10% and 14% of the average welfare generated by the PMA, respectively.

This results from (i) aggregation and (ii) persistent welfare losses at the individual auction level.

Aggregation Firstly, the welfare losses in the RPAs do not all reach their minima at one particular reference price of B, unlike the losses in the SSA which all reached their minima at one particular fixed quantity. The optimal price difference between goods A and B varies across auctions so that the BoE could not commit to a reference price of B that would achieve the welfare-maximising allocation over time. This is the case despite the environment being relatively stable over time.

Auction-level persistent welfare losses Moreover, there are some auctions in which the welfare loss never reaches zero even for the individual auction. This is a more general issue with the RPA, considered in the next section.

Feature 2 The RPA rarely maximises welfare and usually creates significant losses

The welfare losses of the RPAs observed in the period studied is a general feature of the RPA. The reason is that bidders typically submit single bids for large quantities, rather than a set of smaller bids at laddered prices, so even a small change in the reference price of B can cause a large change in the allocation, from a quantity of good B which is much less than the optimal one, to a quantity of good B which is much more than is optimal. Thus even for a given auction, an RPA with a reference price of B close to the optimal price difference for that auction can still generate large welfare losses.[40] Since the optimal difference between

[40]Indeed, even if the fixed price is optimal, there will still in general be welfare losses due to suboptimal rationing of marginal bids. For example, consider the January 2011 auction: even an RPA which turned out, after the fact, to be the best possible one would have generated a welfare loss equivalent to over 2bps of the total quantity allocated that month by the PMA relative to the welfare-maximising outcome. The optimal reference price of B would have been 26bps, but the RPA then accepts all the marginal bids on good B, which represent 29% of the maximum supply (as the auction is uncovered, so the common RPA rationing coefficient equals 1). By contrast, the optimal allocation accepts just 17.2 percentage points of

the clearing prices for goods A and B varies across auctions (this was part of the point of introducing the PMA), any RPA with a price difference that is constant across auctions will lead to significant losses, even when the environment is relatively, but not perfectly, stable.[41]

Feature 3 Welfare losses, BoE revenue and BoE surplus change discontinuously across reference prices but bidder surplus changes smoothly

Figure 4 shows that the differences in net welfare, and BoE revenue and surplus of the RPAs change discontinuously as the reference price of B changes, yet the difference in bidder surplus smoothly increases.

Welfare losses The lumpiness of the welfare loss is driven by the same property of the RPA that causes the mechanism to be suboptimal in general. As the reference price of B changes, the ordering of the bids relative to their reference prices changes. In particular, the ordering of bids on goods A and B changes. If the ordering changes so that a large bid on good B shifts from being allocated on the margin to fully unallocated, the overall quantity allocated of good B decreases significantly, causing a large discrete change in welfare.

BoE surplus and revenue Since the allocation between goods A and B changes in jumps, BoE revenue and BoE surplus also change discontinuously as the reference price of B changes.

Bidder surplus Bidder surplus at the margin is equal to zero for bidders on both goods A and B, because, for each good, the marginal bidder's own bid sets the clearing price. So, a small change in the allocation between goods A and B that does not change the prices does not cause large shifts in bidder surplus. Moreover, the price of A changes very little as the reference price of B changes and often equals zero, so the price for B increases, and bidder

these marginal bids on good B at the spread of 26bps.

For a *given* auction, a fixed price together with a sophisticated rationing rule that mimicked the behaviour of the PMA would, of course, resolve the problem for that auction. However, this would not help much across multiple auctions where different fixed prices would be needed, and where even a slightly "wrong" fixed price for an auction could lead to substantial welfare losses.

[41] Note that this contrasts with the SSA which can perform relatively well in a fairly stable environment if it is correctly calibrated. However, this gives no reason to suppose that either the RPA or SSA would perform better than the other in an environment with significant shocks. (With larger shocks, their relative performance would depend upon the issues discussed in Weitzman (1974) (with price-difference in my context taking the role of price in that paper).)

surplus decreases, almost smoothly with B's reference price.[42,43]

In sum, the impacts of small changes in the RPA reference price are opposite to the results of the SSA, described in Section 4. In the SSA, small changes in the quantities of loans supplied against the two collateral sets cause small changes in the allocation, but can cause large changes in the clearing prices, and hence bidder surplus; in the RPA, small changes in the price difference between the two collateral sets can cause large changes in the allocation, but cause only small changes in the clearing prices and bidder surplus.

Feature 4 Large transfer from BoE to bidders in some RPAs relative to the PMA

In Figure 4, average BoE surplus and revenue are almost uniformly higher in the PMA than the RPA for all reference prices, but bidders as a group only share the welfare gain of the PMA relative to the RPA for high reference prices of B.

When the reference price of B is very small, the market clearing price for good B in the RPA is typically far below the PMA clearing price for B. More bidders win good B in the RPA and, holding the quantity sold constant, bidders benefit more from the RPA than the PMA so that the RPA increases bidder surplus. Conversely, the BoE benefits from the higher revenue in the PMA, and the PMA only allocates bids at prices that the BoE is willing to accept (as measured by the supply curves), whereas the RPA overallocates B relative to the PMA's welfare-maximising quantity, at prices below the supply curve. Because of these two factors, BoE surplus in the PMA is consistently higher than in the RPA.[44]

At the other extreme, when the reference price of B is very large, the difference between bid

[42]Figure 4 illustrates this. When the reference price of good B is zero, bidder surplus is much higher in the RPA than in the PMA. In this case, bids on good A and good B are treated interchangeably and, because the bids on B are almost always higher than bids on A, they are allocated first. Since the cover ratio of good B is always less than 1, all bids on good B are allocated in all auctions. In almost all auctions, the marginal bid is on good A, so that the clearing prices for both goods A and B is around 0-2bps (far below the typical PMA clearing price for good B). This therefore generates a huge bidder surplus in the RPA relative to the PMA. As the reference price of B increases, the ranking of the bids in each RPA changes very little so that the allocation changes very little. The clearing price of B therefore increases almost smoothly with the reference price, causing bidder surplus to smoothly decrease in the RPA.

[43]As the reference price of B increases above 26bps for the 3-month auctions and above 46bps for the 6-month auctions, the bidder surplus in the PMA exceeds bidder surplus in the RPA. This follows from the fact that the marginal winning bid is constrained to be non-negative. When the reference price of B is very large, the bids on good B are below their reference price and therefore go unallocated in the RPA. This leads to a loss of bidder surplus relative to the PMA, in which some bids on good B are typically allocated.

[44]This is analogous to the result in Section 4 that BoE surplus is uniformly higher in the PMA than the SSA. In that case, the difference acts through the price determination rule — the PMA clearing prices depend on the relative supply curve, whereas the SSA clearing prices do not.

and reference price for many bids on good B is negative, so that they are unallocated. The RPA therefore underallocates B relative to the PMA, which allocates the efficient quantities, and raises the price paid for B, so that bidder surplus is lower and BoE surplus is higher in the RPA than in the PMA.

As for the SSA, this transfer from BoE to bidders is partly driven by the price determination rules, as the PMA moves from an LWP rule to an HLP rule. However, I calculated the average differences in surpluses for the PMA with an HLP rule for the entire period and the RPA and found only a small change in transfers.[45]

Feature 5 Fixed prices could also cause large welfare losses

A simpler alternative mechanism would be for the BoE to fix two prices, one for each of goods A and B, at which bidders could borrow funds. However, this would likely cause significant welfare losses, in excess of those estimated for the RPA over this period for two reasons. Most importantly and akin to the difficulty of choosing the reference price of B in the RPA, the BoE would have no ability to predict the optimal fixed prices prior to the auction, so that the posted prices would possibly be far from optimal. Secondly, if the total quantity of funds was capped and excess demand was rationed, the loans would not go to the bidders that value it most, causing welfare losses.[46]

To illustrate this, I compared the surpluses for the PMA to posting a pair of fixed prices, at a range of different levels, with a fixed supply, set equal to the BoE's maximum supply, and winning bids rationed pro-rata. Figures B.1 and B.2 show these differences in surpluses as the fixed price for Level B changes, holding constant the fixed price of Level A, averaged over the 3-month and 6-month term auctions, respectively. These are qualitatively similar to the differences in surpluses between the PMA and RPA. I would expect the differences between the PMA and fixed prices to be quantitatively larger in a period of greater stress.

[45]There are some minor differences. For example, the benefit to bidders of the PMA is small (but positive) when comparing an RPA with a large reference price of B to the PMA-HLP, but is zero when comparing with the actual version of the PMA used.

[46]If, instead, the total supply of funds was uncapped, posting fixed prices may lead to an extremely adverse welfare outcome, depending on the prices that the BoE would be willing to accept to supply funds beyond the BoE's stated maximum supply. For example, if the prices that the BoE was willing to accept for goods A and B were constant beyond the stated maximum supply and equal to the prices at that maximum supply, there might be a large adverse welfare outcome, especially if the BoE underestimates the demand for Level B at the posted price.

6 Welfare and surpluses in periods of stress

Demand was relatively stable in the period studied, so that the welfare benefits, arising from the flexibility of the allocation in the PMA, were fairly muted. To investigate the PMA's relative performance in a period of greater stress, I compare the surpluses and welfare of the PMA with the SSA and RPA given simulated bids. I transform the observed bids by multiplying the observed quantities demanded by 4 and 8 for the bids on Levels A and B, respectively, and multiplying the observed bid prices submitted for Level B by 4. This is intended to loosely reflect the pattern of demand at the beginning of the COVID-19 pandemic in March 2020. Figures B.3 and B.4 repeat the calculations of the differences in surpluses using these simulated bids.

The welfare gain of the PMA is qualitatively similar, but quantitatively significantly larger in absolute terms, in this exercise simulating stress, reflecting the scaling up of benefits arising from the flexibility of its allocation across collateral types.

As in the main analysis, the BoE's surplus is always larger—and is now very much larger—in the PMA than in either the SSA or the RPA in this experiment. Importantly, because the PMA's allocations respond to the BoE's supply preferences, the BoE obtains (weakly) positive surplus from every unit it allocates, which is not always the case in the SSA and RPA, whose allocations are insensitive to its preferences.

However, the impact on bidders' surplus differs across comparators: the bidders gain from the PMA over a larger range of SSAs, but over a smaller range of RPAs.

Figure B.3 shows that the PMA now gives bidders more surplus (in aggregate) than all SSAs that allocate less than 45% to Level B for 3-month terms, and all that allocate less than 75% for 6-month terms. For the original demands, the PMA only gave bidders more surplus than SSAs that allocated less than 10% and 20% to Level B, respectively. The reason for these results is that the relatively flat demand for loans secured by Level A (which is preserved in the transformations) means bidders can never gain much surplus on Level A. So when there are a large number of high bids on Level B (as would be expected during a period of stress), bidders gain more surplus from the PMA, which will allocate a large share of the maximum supply to B in these circumstances, than from an SSA, unless the SSA also allocates a large share to Level B.

By contrast, Figure B.4 shows that the PMA now only gives bidders more surplus than RPAs with a reference price of B higher than 78bps for 3-month terms and higher than 68bps for

the 6-month terms. For the original demands, the PMA gave bidders more surplus than RPAs with a reference price of B higher than 26bps and 48bps, respectively. The reason is that in periods of greater stress bids on B increase relative to A, so for any reference price of B, the RPA accepts more on Level B, generating larger net surplus to the winning bidders. The PMA also accepts some more on Level B relative to the original demands, because it is more efficient to do so when bids on B are relatively high. But the change in the RPA is typically larger—to the benefit of the bidders—as the PMA also takes the BoE's supply preferences into account. So, under stressed conditions, bidders prefer a PMA to an RPA only if the RPA has a very high reference price of B.

I experimented with various alternative transformations of observed demand and they yielded broadly similar results. However, the results about the bidders' preferences between the designs may not generalise beyond the specific case of two eligible collaterals of which one is Level A, because these results seem to be caused by the flat demand on Level A.

7 Alternative objective functions

The PMA used by the BoE in the period studied is designed to maximise total welfare. I find that changing this objective to maximising BoE surplus has limited qualitative impact on the results. Figures B.5 and B.6 compare the surpluses of the SSA and RPA to a PMA, in which the allocation adjusts to maximise BoE surplus.

The welfare gain of the PMA is broadly similar both qualitatively and quantitatively to the case in which the PMA maximises total welfare. This reflects the importance of the PMA's responsiveness to the BoE's supply preferences in driving the welfare gains relative to the SSA and RPA.

The distributional impact, that is, the impact on the bidders' and BoE's respective surpluses, is qualitatively similar but quantitatively larger to the main analysis. This is unsurprising as the BoE-surplus-maximising PMA systematically favours the BoE relative to the bidders, amplifying the observed benefits to the BoE of the PMA over the sample period.

The comparisons for a PMA that maximises BoE revenue are broadly similar. (The welfare gain of a PMA which maximises BoE revenue, i.e. the total payment from bidders to the BoE, is slightly smaller than that of the BoE-surplus-maximising PMA, whereas the transfer of surplus from bidders to the BoE is also typically slightly smaller.)

8 Broader benefits of the PMA

Unlike the SSA and RPA, the PMA adjusts both the price and the quantity allocated of each good in response to the bidders' and BoE's preferences (expressed by the bids and the supply curves, respectively). The benefits of this automatic adjustment can be split into the direct welfare gains of meeting the liquidity needs of individual participants, and the externalities that Fisher (2011a) describes of providing liquidity insurance to the system as a whole, avoiding moral hazard, and informing the BoE about stress in the market (in addition to communicating the "correct" price to the market).[47]

One measure of the PMA's performance in terms of the overall benefits would be how far the outcome of an alternative design would be from the optimal prices and quantities. To the extent that the BoE's supply curves express the optimal price-quantity relationship, the PMA is designed to allocate the "correct" quantity at the "correct" price but the alternatives do not necessarily do so.[48] The alternatives' allocations would not just be inefficient as measured by the spreads on the loans, but also for society more broadly. How one measures this is a matter of choice — it could be the price difference from the optimum, the quantity difference, or the product of the two. My measure is roughly proportional to the product of the two.[49]

The difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept on the loans allocated in the auctions does not capture all the benefits of the PMA. My measure is therefore a conservative estimate of the overall benefits.

[47] Only some of these broader benefits are relevant to the comparison of the PMA to each alternative. The benefit of providing system-wide liquidity insurance is particularly important in the comparison of the PMA to an alternative which supplies too little funds against Level B, but less important when comparing the optimum to an alternative which supplies too much. The benefit of avoiding moral hazard is more important in the converse case where too much is supplied against Level B.

[48] The PMA allocates the quantities at which the BoE's supply curves meets the bidders' demand curves and it sets prices to implement this allocation. In contrast, in the SSA for example the allocation of funds is fixed for each collateral set so, in a period of stress, demand against Level B might be far higher than the BoE's supply curve at this fixed allocation meaning the allocation of B would be far lower than optimal.

[49] It would be exactly equal if the BoE's supply curve and the bidders' demand curve were linear with equal slopes in absolute value for each good. It is smaller than the product of the two if the average slopes of the supply curves are smaller than the average slopes of the corresponding demand curves in absolute value between the quantities sold in the two designs, the supply curves are weakly convex—which approximates the BoE's supply curves providing the quantity allocated against Level B is not too large—and the demand curves are weakly concave. It is larger than the product of the two under converse conditions.

9 Conclusion

I have shown that the flexibility in the allocation of the PMA generates welfare gains relative to the two simple alternatives of the SSA and RPA. The RPA could not have maximised welfare with any fixed reference prices, and hence fixed price difference, due to the instability of the optimal price difference and the lumpiness of bids. And although the SSA could have almost maximised welfare in the period had the BoE been able to accurately forecast the optimal quantities, this result was specific to the stability of the optimal quantities in this particular time period. Moreover, it seems unrealistic to think that the BoE could have forecast these quantities.

The BoE always gained from the PMA relative to the comparators in the period studied. However, the welfare gains of the PMA were not always shared with the bidders. Average bidder surplus was higher in the PMA relative to the comparators only when either the fixed share to Level B used in the SSA was low, or when the fixed reference price of B in the RPA was high. Three features in the period can explain why the bidders did not always gain: the absence of significant stress, only one collateral whose demand is not flat, and little paired bidding. Absent any one of these three features, the bidders, as well as the BoE, might have benefited more from the use of the PMA. Indeed, bidders might benefit more under the current ILTR design relative to the sample period. Since 2014, participants can bid for loans secured by an additional collateral set, which is less liquid than Levels A and B and for which I would not expect demand to be flat. Moreover, the design might have been more beneficial to bidders at the beginning of the COVID-19 pandemic in March 2020, in which demand for loans was acute.

However, the analysis has ignored several significant advantages of the PMA.

One benefit of the PMA relative to the alternatives is that the PMA reduces the scope for strategic behaviour.[50] The PMA combines the markets into a single mechanism, encouraging competition and reducing bidders' incentives relative to the SSA to unilaterally understate their demands. It also may reduce the bidders' ability to coordinate prices relative to the alternative designs, especially when the auctioneer's preferences are not revealed to the bidders, as in the ILTRs. I maintained an assumption of truthful bidding across all designs, eliminating this potential source of welfare loss. If this fails to hold, the actual welfare losses of the SSA and RPA relative to the PMA may in fact be much larger.

[50] Chapter 1 tests for strategic behaviour in the ILTR operations and finds that truthful bidding is a reasonable approximation under the PMA in the ILTR auctions.

Furthermore, even in the absence of these effects, the welfare gain of the PMA would likely have been much greater in a period with less stable demand.

Most importantly, neither the SSA nor the RPA would have achieved the BoE's main objectives in setting up the PMA. Any SSA is simply setting a quantity for funds lent against the less liquid Level B collateral, and any RPA is simply setting a fixed spread between loans secured by collateral assets with different liquidity levels. So neither of these natural alternative mechanisms to the PMA would both allow a substantial increase in allocations of funds against Level B collateral in the event of stressed conditions, but only allow this at higher interest rate spreads, as the BoE wished. This chapter has shown that the PMA enhanced efficiency as measured by the difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept on the loans allocated in the auctions. Broader benefits of the auctions, e.g. the ability to stop runs on individual banks (and the externalities of avoiding systemic failure) while avoiding moral hazard, are taken into account, to the extent that the BoE's supply curve internalises some externalities. Paul Fisher (then Executive Director at the BoE) described the design as "potentially a major step forward in practical policies to support financial stability" (Milnes, 2010). The automatic adjustment to market conditions of both the amount of funds loaned in the PMA, and the interest rate premium charged, improves efficiency not only directly, as analysed in this chapter, but also in this broader sense, too.

Appendices

A The Bank of England's algorithms

A.1 Detailed description

I simplify by ignoring the exceptional knife-edge cases, that never actually arose, in which the competitive equilibrium prices were not unique, so write the competitive equilibrium prices and quantities as $p_A^c, p_B^c, q_A^c, q_B^c$, respectively.[51] Based on the results, and on available information about how the algorithm changed in October 2011 and March 2013:[52]

The BoE's programs always set $q_A = q_A^c$.

(a) From the beginning until September 2011: The BoE set p_A equal to the highest price for which demand on A equalled q_A^c. It then computed q_B based on p_A and the relative supply curve. It then set p_B equal to the highest price for B that was consistent with a demand for q_B (which price could exceed the unique price consistent with a supply of q_B).

(b) From October 2011 until February 2013: The BoE set p_A and q_B as in (a) above. It then set p_B equal to the unique price for B that was consistent with both the demand for, and the supply of, q_B.

(c) From March 2013 to January 2014: The BoE chose the competitive equilibrium prices and quantities.

[51] If the auction is not fully allocated (in competitive equilibrium), p_A^c is unique ($p_A^c = 0$), and the monotonic strictly upward-sloping relative supply curve then uniquely determines p_B^c (at the intersection of the relative supply and the demand curve for good B). If, instead, the auction is fully allocated, p_A^c is uniquely determined except in the knife-edge case in which the numbers of bids taken on A and B exactly "fit" the amount to be allocated without any need for rationing on good A. Even in this case, the difference between the prices, $p_B - p_A$, is uniquely determined, since the relative supply curve is upward sloping (see fig. 2 of Klemperer, 2008, for illustration). Moreover, q_B^c is therefore also uniquely determined. So only if there is also no need for rationing good B, is there any ambiguity in the prices. (Similarly, in a simple uniform price auction, only if the exact quantity offered is sold without rationing is there a distinction between "HLP" and "LWP".)

[52] There is less evidence for the period March 2013 – January 2014 than for the other periods, since there were no bids on good A at prices above 0 prior in this period. However, the BoE definitely intended $p_A = 0$ for non-fully allocated auctions starting in 3/2013, and competitive equilibrium prices and quantities should have been the result. A completely different algorithm was used to find competitive equilibrium prices and quantities from February 2014.

These changes in the PMA price determination rules are summarised in Table 1 in Section 2.

A.2 Implications for welfare

In all fully allocated auctions, and in non-fully allocated auctions in which $p_A = 0$, the results were always as in Klemperer (2008), except in the first period, period (a) above, in which sometimes $p_B > p_B^c$.[53] Since the quantities were the same (even in the first period) as in the originally-defined PMA, i.e. $q_A = q_A^C$, and $q_B = q_B^c$, it follows that the welfare (as conventionally defined) was also the same.

In the two auctions in the first period, (a), and the five auctions in the second period, (b), which were non-fully allocated with $p_A > 0$, the results were as in Klemperer (2008) except that $p_A > p_A^c$ (since $p_A^c = 0$) and sometimes $p_B > p_B^c$ and sometimes $q_B < q_B^c$.[54] However, only in January 2011, December 2011, April 2012, and June 2012,[55] was $q_B < q_B^c$, so only in these auctions were the quantities allocated different from Klemperer (2008) and only on good B. Only in these four auctions, therefore, did social welfare (as conventionally measured) differ from Klemperer (2008).

A.3 Interpreting the BoE's preferences

The PMA developed for the BoE is described in Klemperer (2008). The proposal was to implement competitive equilibrium allocations and prices, assuming that the BoE's relative supply curve and the counterparties' bids represent their true preferences.[56] (In the event of non-uniqueness, Klemperer (2008) recommended choosing the lowest competitive equilibrium prices,[57] so that small bidders, at least, have no incentive to misrepresent their valuations.[58])

[53] p_A^c was automatically identical to the good A clearing price in the originally defined PMA in fully allocated auctions, except in exceptional "knife-edge" cases that never arose.

[54] It is theoretically possible, though it never happened, that an auction could have been fully allocated in competitive equilibrium but not fully allocated by the BoE in periods (a) or (b). (The reverse could not happen.)

[55] The three auctions of these seven in which $q_B = q_B^c$ were April 2011, March 2012, and October 2012.

[56] This was justified by the expectation that no bidder could gain much by misrepresenting her preferences.

[57] The lowest competitive equilibrium price vector is uniquely defined—it is the most "south-westerly" point on a 45° line in fig. 1b of Klemperer (2008), if the competitive equilibrium price vector is not unique.

[58] If a bidder expects LWP, she has an incentive to misrepresent her preferences, regardless of her size. She can increase her expected surplus by reducing her bid below her value, which lowers the clearing price, increasing her surplus, in the case in which her bid determines the clearing price. (Large bidders have the additional incentive to misrepresent her preferences by exerting market power, under both the HLP and

The PMA initially implemented by the BoE was a close approximation to, but differed from, the PMA as originally proposed. It sometimes yielded different results, which create some ambiguity in the appropriate measure of welfare. (Appendix A.4 shows that these differences were infrequent and small, and the BoE subsequently moved in two steps towards an exact implementation of Klemperer (2008).)

If the BoE's relative supply curve represents the minimum price that the BoE is willing to accept to supply good B relative to good A (as a function of the quantity of good B allocated), then the social welfare from any given auction outcome is unambiguous and easy to calculate.

The outcomes of the initial implementation and originally defined PMA differed because choosing competitive equilibrium prices is not always the same as first choosing competitive equilibrium allocations and then choosing (either lowest or highest) prices consistent with total bidder demands equalling those allocations.[59] This is because the auctioneer might prefer different (either higher or lower) allocations at those prices. As discussed in Appendix A.1, above, this difference only arose when cover was low enough that an auction was not fully allocated.

To see this, suppose the only bids are a small bid on good A at 10bps, and a small bid on good B at 12bps, and the BoE's relative supply (of good B) for small allocations is 5bps. Then the only competitive equilibrium prices are $p_A^c = 0$bps and $p_B^c = 5$bps, reflecting the minimum prices that the BoE is willing to accept. With any higher prices on either good, the BoE would not be in equilibrium, since it would wish to sell more than was demanded. A result of this kind arose in the June 2012 auction, in which the BoE's programme chose $p_A = 10$bps and no sale to good B (as if $p_B = 15$bps).

A revealed-preference approach would be that the BoE was in fact always optimising its welfare function. The BoE's supply function would indicate a preference for a price difference of 5bps in this case, so that if $p_A = 10$bps—which would allow it to serve all the demand on good A—then the correct price for p_B should be 15bps. The problem is that it is then hard to know exactly what the BoE's welfare function was. It is unclear why $p_A = 10$bps,

LWP.)

[59] The BoE's initial approach was close to choosing competitive equilibrium allocations, and then choosing the highest prices consistent with total bidder demands equalling those allocations (and was always equivalent to this for fully allocated auctions). So for fully allocated auctions the BoE's initial approach corresponded to lexicographically maximising social welfare as conventionally defined, then, subject to this, maximising BoE surplus. (Klemperer (2008) discouraged this choice of higher than competitive prices, since it gave bidders incentives, albeit probably small ones, to misrepresent their preferences. Hence the BoE's moves to their subsequent implementations.)

$p_B = 15$bps is superior to $p_A = 0$bps, $p_B = 5$bps, or other choices with $p_B = p_A + 5$. For example, $p_A = 7$bps, $p_B = 12$bps are the highest prices consistent with total bidder demands equalling those allocations, and might be more "profitable".[60] And $p_A = 0$bps, $p_B = 5$bps yield the highest social surplus if I assume the "disutility" (i.e. "cost") to the BoE of allocating good B is $p_A + 5$. Even if there is a clear reason why $p_A = 10$bps, $p_B = 15$bps is preferred,[61] it would be unclear what assumption should be made about the BoE's utility function when making comparisons with the RPA and SSA which won't generally yield the "correct" price difference.[62]

I have therefore assumed that the BoE all along intended to implement the PMA as originally defined, as is consistent with its behaviour in the latter part of the period studied. That is, I assume that the BoE's objective throughout was to maximise the sum of bidder surplus and BoE surplus, as defined in Section 3.3.

A.4 Relevance of the HLP versus LWP choice to measured welfare

Welfare, as defined in Section 3.3, is only affected by the difference between the BoE's implementation and Klemperer (2008) when all bids on good A are allocated, and the lowest bid on good A is strictly positive, so that the clearing price for good A differs under the two pricing rules. Under LWP, it is strictly positive and under HLP, it equals zero. In this case the allocation of B may differ between the two rules (in which case, at least under LWP, B will not be fully allocated).

Assuming, as in Appendix A.3, that the BoE's relative supply curve represents the minimum price it is willing to accept to supply good B as a function of its quantity, the BoE's actual

[60] Choosing $p_A = 7$bps, $p_B = 12$bps would have yielded a greater BoE surplus (assuming the prices that the BoE was willing to accept were 0bps and 5bps) in June 2012.

[61] Perhaps the BoE wanted the highest price for good A that was consistent with demand matching the competitive equilibrium quantity. And perhaps it also had a strong preference for the "correct" spread. The BoE likely did have preferences over the price-spread (independent of the prices that the BoE was willing to accept, as revealed by the supply curves), in order to send an informative signal to the market, and to charge an appropriate premium for loans secured by poor-quality collateral. If, in our example, the small bid on B was at 5bps rather than 12bps, and if the BoE wanted $p_A = 10$bps, it would have been odd to allocate the bid on B, which would have required $p_B < p_A$ but have yielded no welfare gain (based upon the bids and supply curves).

[62] For example, how much worse for the BoE are $p_A = 10$bps, $p_B = 12$bps, which are the highest prices consistent with total bidder demands equalling the competitive equilibrium quantities? To what extent do the extra allocation of B (and extra BoE surplus if one assumes the prices the BoE is willing to accept are 0bps and 5bps) mitigate the disutility of the price difference being smaller than the BoE's relative supply? Note that prior to the October 2011 revision of the algorithm the BoE sometimes chose price differences larger than its relative supply, though it never chose smaller price differences.

Table .1: Auctions in which the actual PMA is possibly suboptimal in the ILTR auctions, June 2010 – January 2014

Operation date	Amount allocated in actual PMA (% of auction)		Optimal amount allocated (% of auction)		Cover (% of auction)		Clearing spread (bps)	
	Good A	Good B	Good A	Good B	Good A	Good B	Good A	Good B
11 January 2011	54.4	15.65	54.4	17.6	54.4	40	5	26
13 December 2011	16.6	9.85	16.6	11.5	16.6	14.36	3	12
10 April 2012	4	2.7	4	3.7	4	5	4	10
12 June 2012	0.1	0	0.1	2.8	0.1	2.8	10	N/A

implementations were slightly suboptimal on a few occasions. Table .1 shows the actual and optimal amounts allocated, cover ratios and clearing prices in the four auctions in which the allocation in the sample period differs from the competitive equilibrium allocation.

This issue is not economically significant. It affects only four low-cover auctions in the period studied—January 2011, December 2011, April 2012, and June 2012—and the resulting measured welfare loss is very small. The difference in net welfare between the SSA and PMA is 0.06bps, 0.06bps and 0.15bps of the actual PMA allocation in these auctions, respectively. Moreover, as discussed above, it is not clear whether this measured welfare loss is real, since there is some ambiguity about the BoE's preferences in these cases.

A further question then arises as to whether to compare the RPA and SSA with what the BoE actually did, or to what the outcomes would have been if the BoE's final implementation (i.e. the PMA-HLP) had been used throughout. The latter choice is perhaps the most natural one, since it is the most relevant guide for future policy. In the main analysis, consistent with the other assumptions I have made, I have taken the conservative approach that minimises the average welfare gain of the PMA over the RPA and SSA, and compared the RPA and SSA with the outcomes of the BoE's actual implementations. That is, I have assumed the BoE made slight mistakes in some early auctions.

Because they were only in a minority of low-cover auctions, they are relatively unimportant and have limited impact on the results. I repeat the surplus difference calculations using HLP for the PMA price determination in all auctions, denoted PMA-HLP. Figure B.7 shows the differences in auction surpluses between the PMA-HLP and SSA, averaged over the June 2010 – January 2014 auctions, weighted by the quantities allocated in the PMA. Figure B.8 shows the analogous surplus differences between the PMA-HLP and RPA.

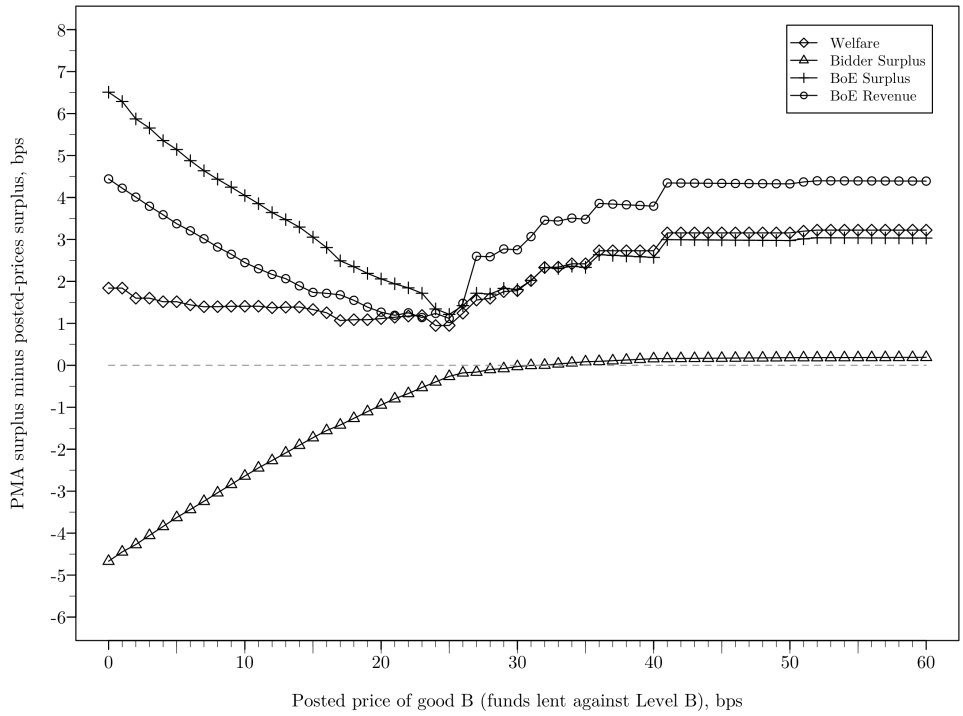
The net welfare differences are almost identical to the corresponding results for the actual PMA, shown in Figures 3 and 4. The transfers from bidders to the BoE under the PMA are slightly smaller, as the clearing prices are on average lower when using HLP.

Since the PMA price determination rule has been HLP on both goods since March 2013, the ambiguity or potential welfare loss can no longer arise.

B Additional figures

Figure B.1: Difference in auction surpluses, PMA minus posted fixed prices with Level A price fixed at 0bps, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

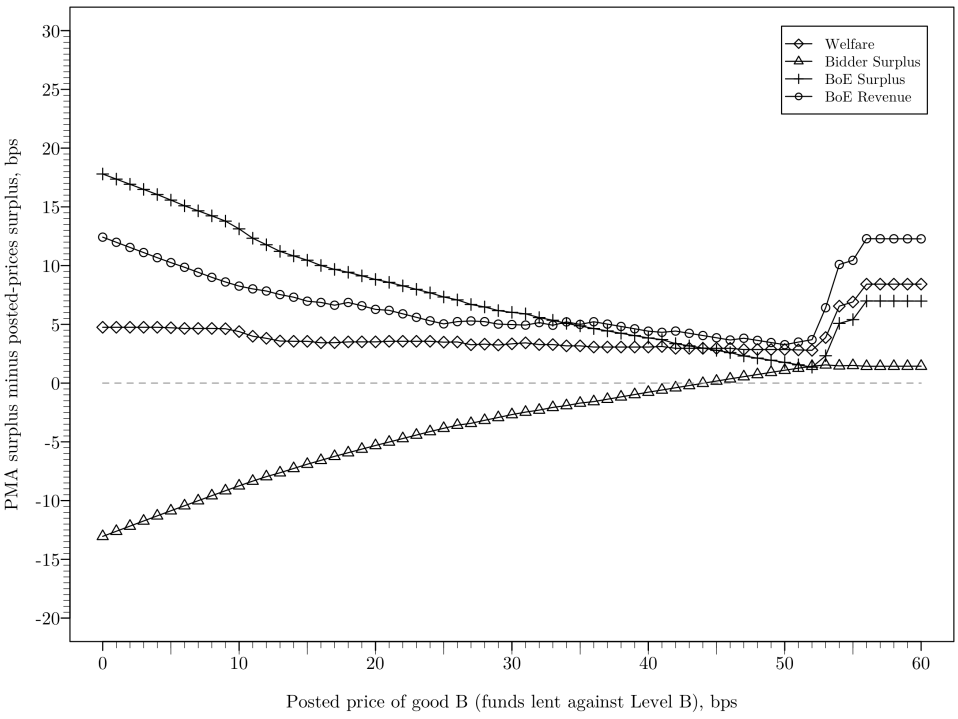
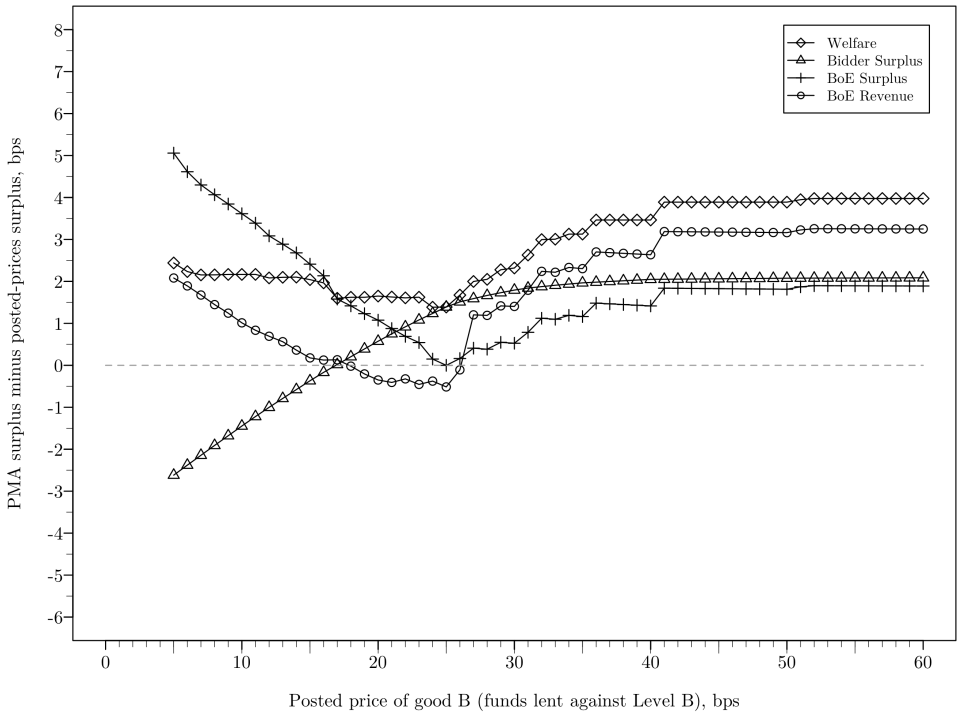


Figure B.2: Difference in auction surpluses, PMA minus posted fixed prices Level A price fixed at 5bps, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

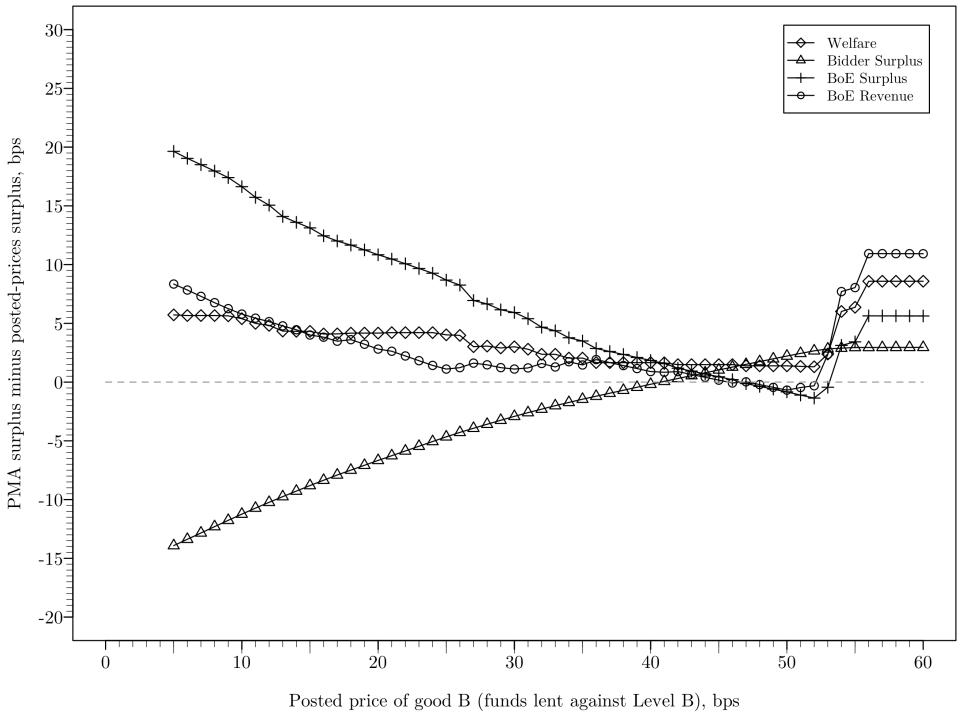
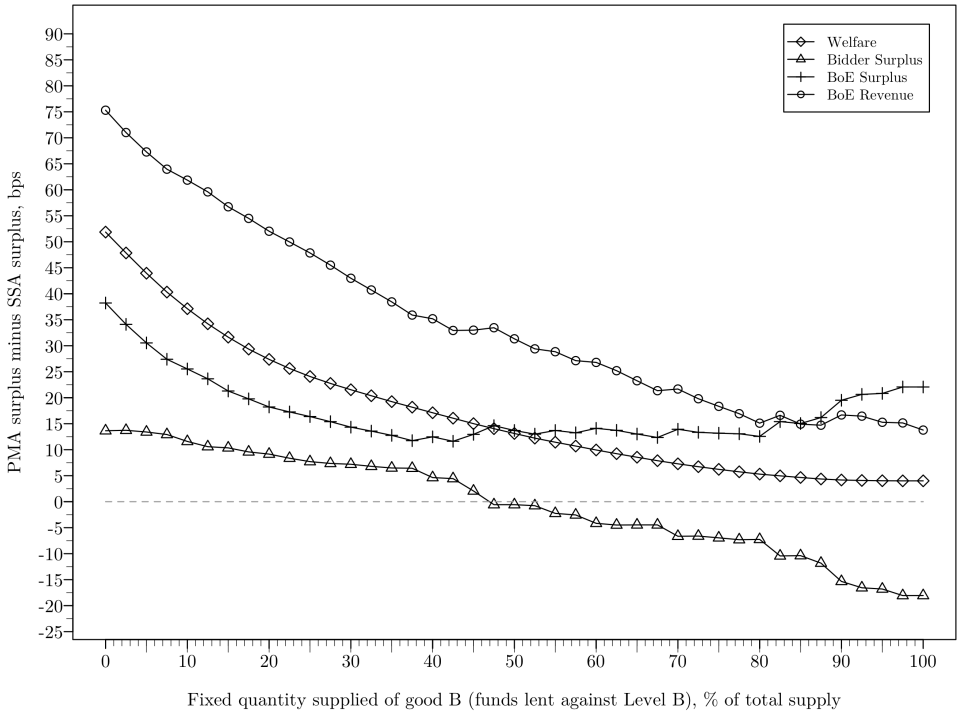


Figure B.3: Simulated difference in auction surpluses, PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation), transformed bids: Level B quantities demanded multiplied by 8, Level A quantities demanded multiplied by 4, Level B bid prices multiplied by 4

(a) 3-month term



(b) 6-month term

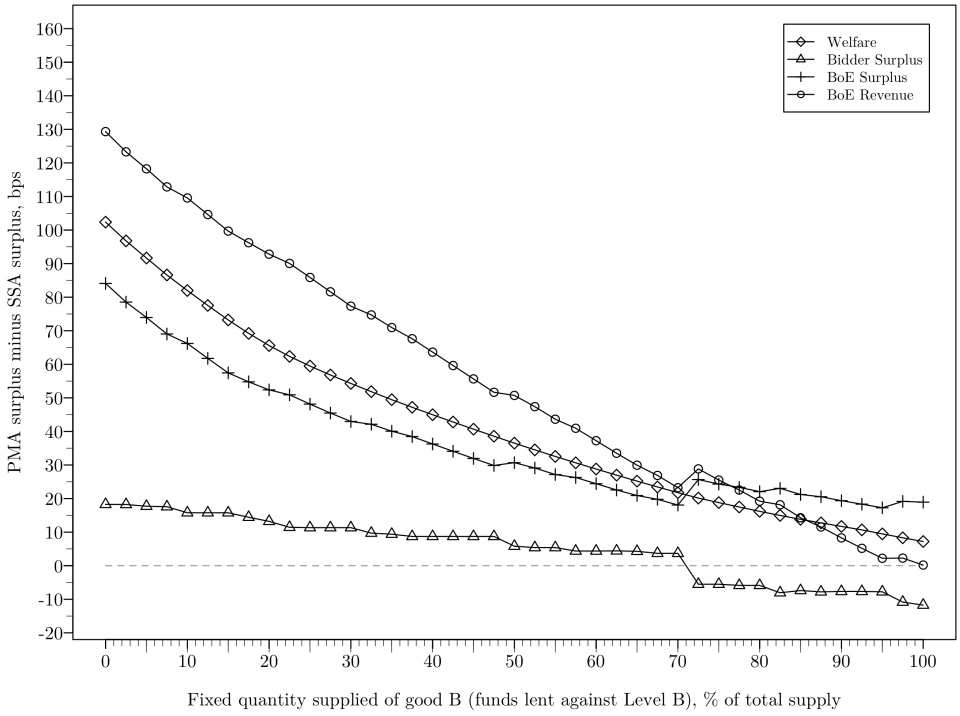
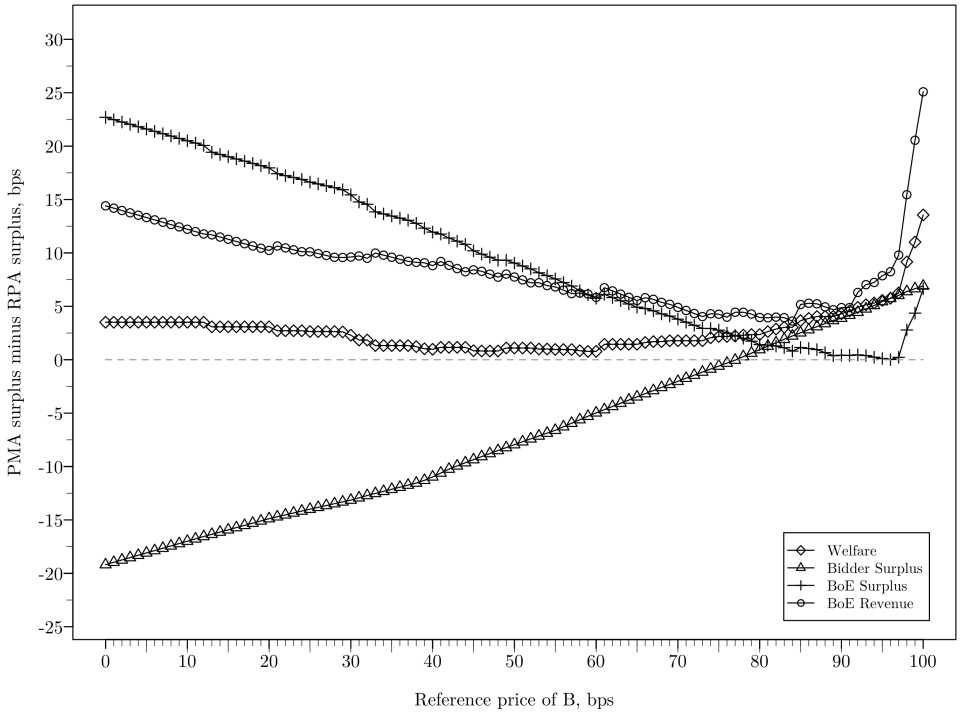


Figure B.4: Simulated difference in auction surpluses, PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation), transformed bids: Level B quantities demanded multiplied by 8, Level A quantities demanded multiplied by 4, Level B bid prices multiplied by 4

(a) 3-month term



(b) 6-month term

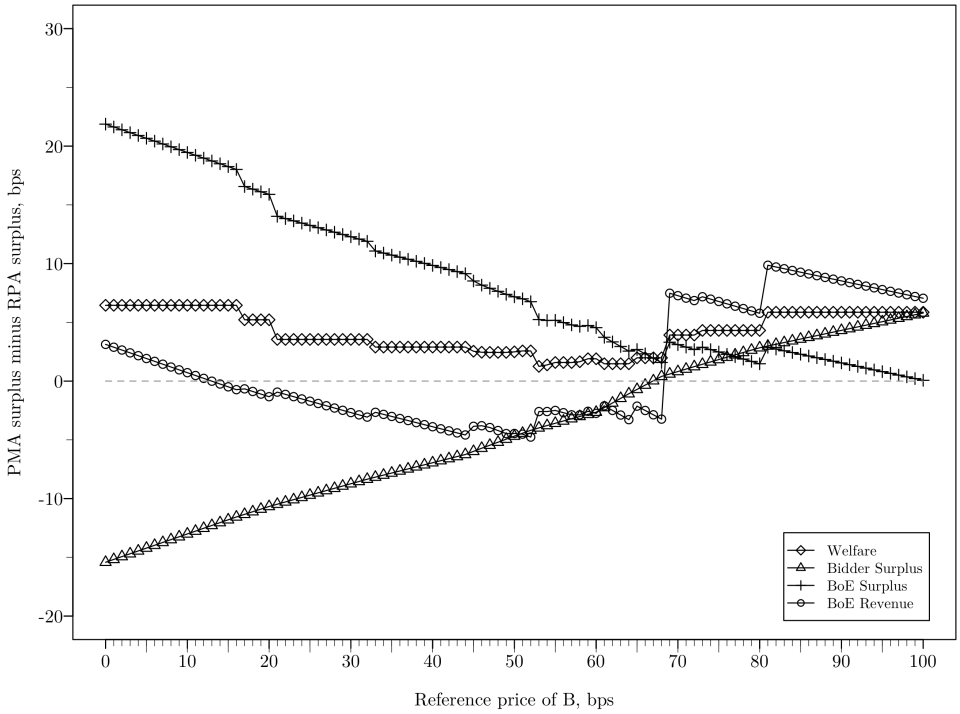
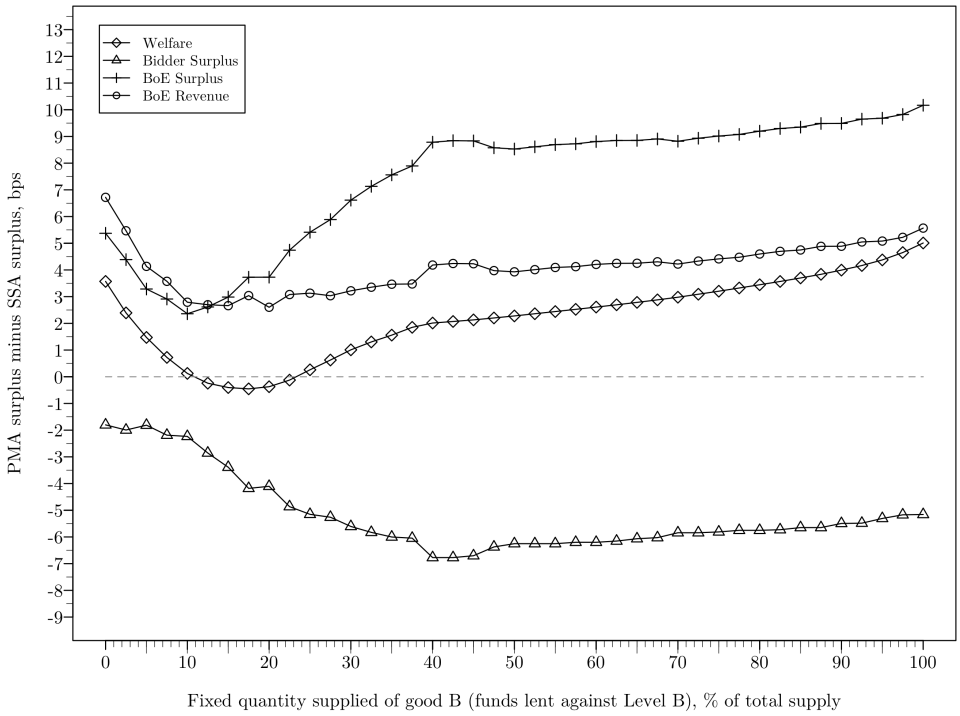


Figure B.5: Difference in auction surpluses, BoE-surplus-maximising PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

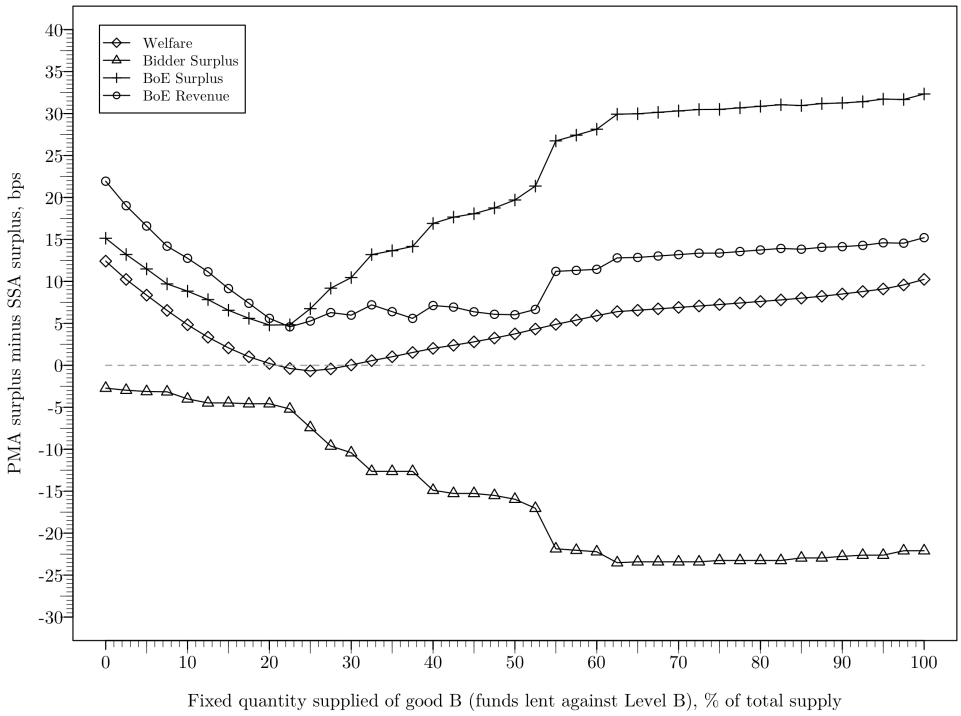
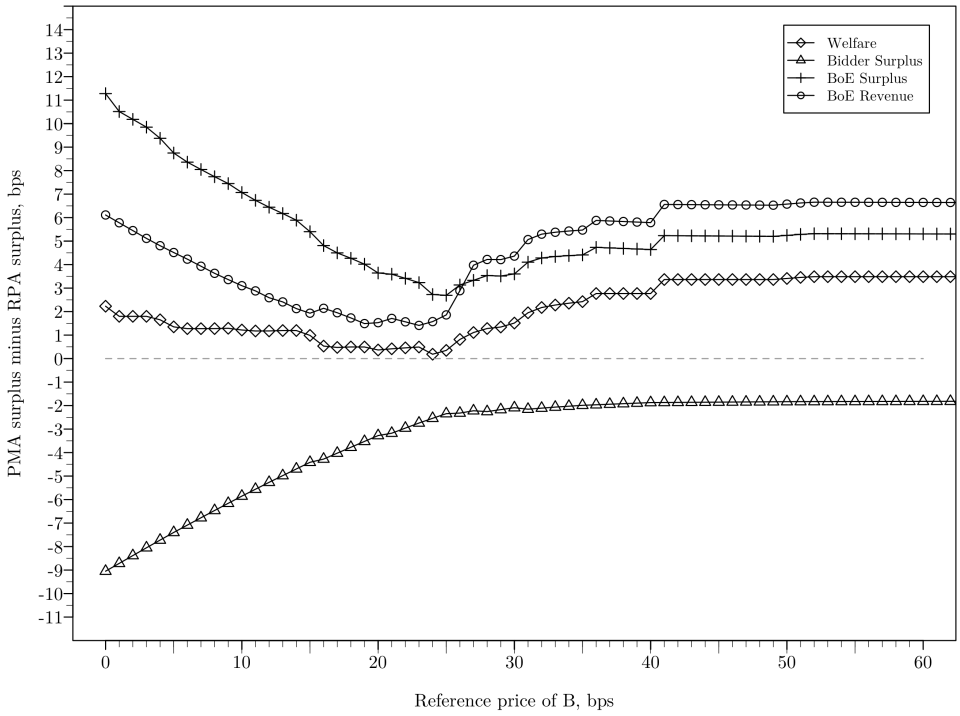


Figure B.6: Difference in auction surpluses, BoE-surplus-maximising PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

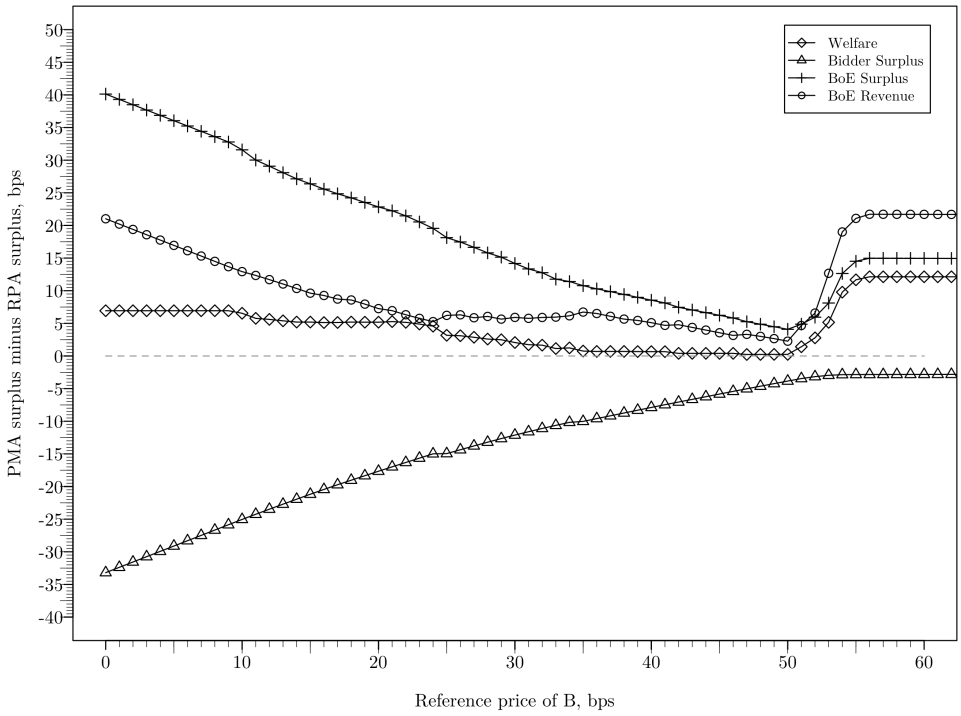
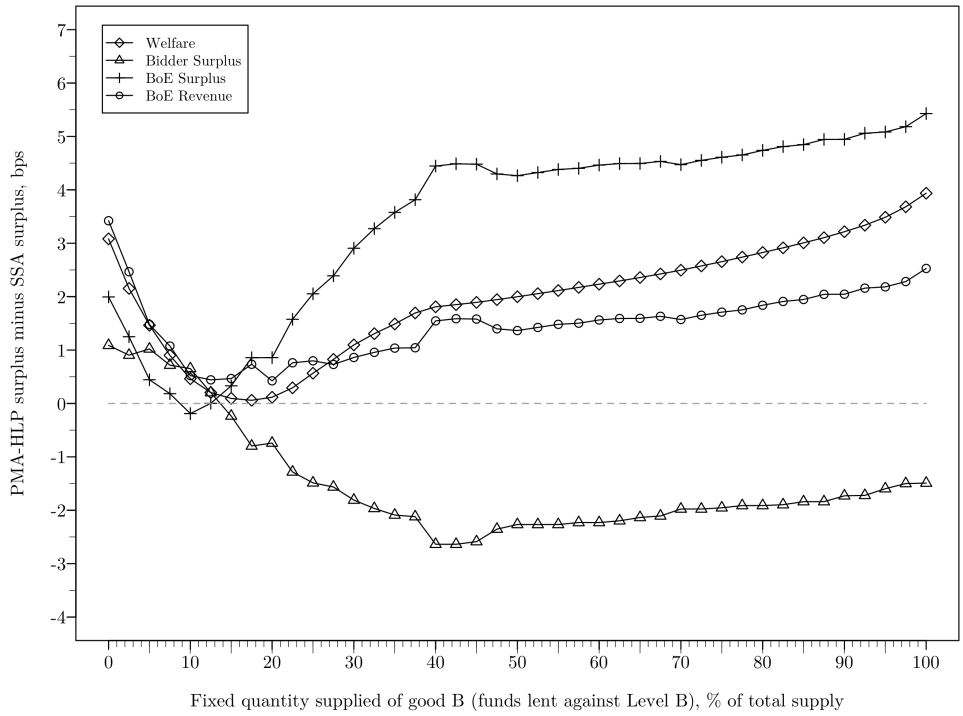


Figure B.7: Difference in auction surpluses, PMA-HLP minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

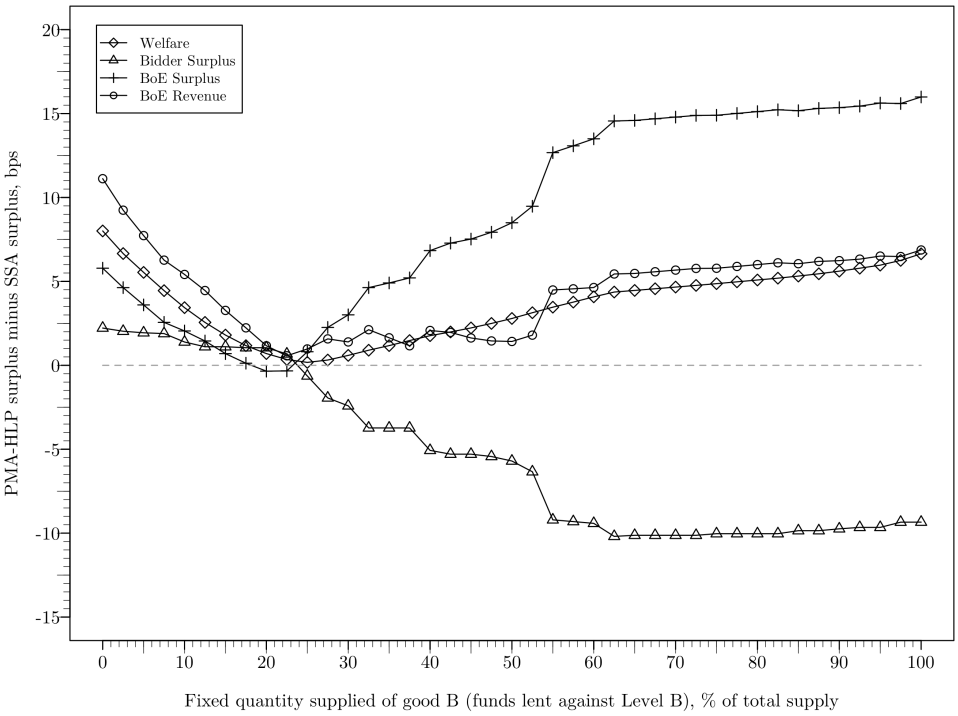
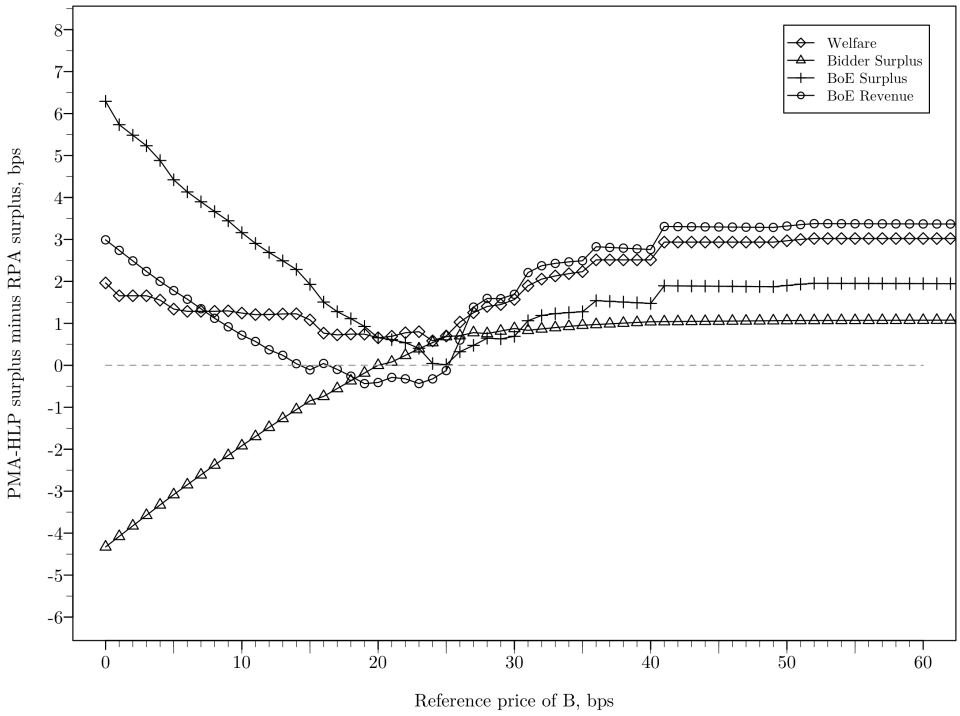
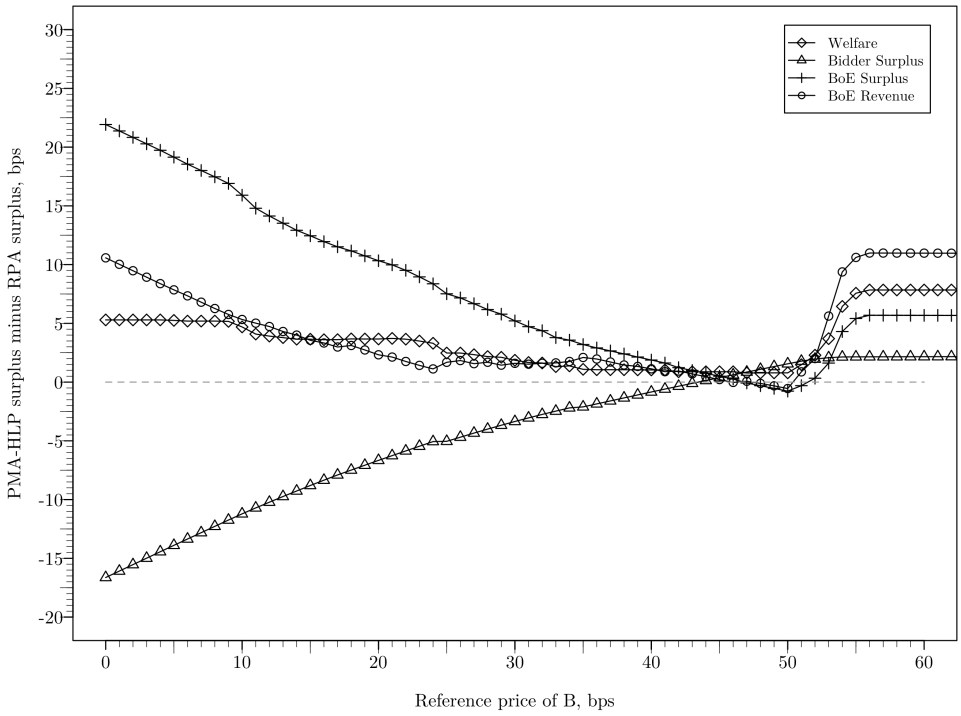


Figure B.8: Difference in auction surpluses, PMA-HLP minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term



CHAPTER 4
Who gains from efficient auctions?

Abstract

I provide conditions under which increasing efficiency by combining competitive auctions of substitute goods increases *both* expected bidder and auctioneer surplus relative to running separate simultaneous auctions. I also provide conditions under which one side's expected surplus falls. I show how the distributional effects depend on the bidders' aggregate bid functions and the auctioneer's costs of production. I estimate the magnitude of these effects using novel bid-level data from Mexico's domestic government bond market.

1 Introduction

Multi-unit auctions are often held separately and simultaneously for different goods. For example, in the primary market for government debt, independent auctions are held for bonds of different characteristics by the same institution at the same time.

This is inefficient if either the auctioneer or the bidders view the goods as substitutes. Selling goods using *multi-product* multi-unit auctions can increase efficiency. In these auctions, the allocation across goods depends on the submitted bids and the auctioneer's own supply preferences. They are straightforward to implement and have been used in both financial markets and renewable energy procurement and support.[1]

Switching to the optimal multi-product auction naturally increases efficiency if participants bid competitively for substitute goods.[2] The effect on the distribution of surplus between bidders and the auctioneer is less clear. Who gains, and who loses? I provide conditions under which the optimal mechanism benefits both sides of the market. I also provide conditions under which one side of the market, either the bidders or the auctioneer, disproportionally

[1] In financial markets, a notable use of multi-product auctions is quantitative easing: both the Federal Reserve and the Bank of England held multi-product auctions, in which the quantities purchased of different government bonds were a function of the submitted bids and the central banks' willingnesses to pay (expressed by the bonds' secondary market prices). Within these auctions, the price differences between bonds were fixed. This mechanism is efficient when the marginal benefit to the central bank of one bond relative to another is constant (see Section 7.3). The Bank of England runs other multi-product auctions, too. Its Corporate Bond Purchase Scheme—another component of quantitative easing—is designed to jointly purchase multiple corporate bonds. Its weekly Indexed Long-Term Repo operations are Product-Mix Auctions, which jointly allocate liquidity to financial institutions against multiple types of collateral. Other examples, which motivates my empirical analysis in Section 7, are from Mexico. Banco de México, the Mexican central bank, conducts swaps of a range of government securities (differing in, e.g., their term, coupons, and whether nominal or real interest rates) within multi-product auctions called "subastas de vasos comunicantes". It also conducts similar swaps, referred to as "permutas de valores gubernamentales", on behalf of the Mexican federal government.

In renewable energy procurement and support, multi-product auctions are increasingly popular to allocate contracts across different technologies (Szabó et al., 2020). Introduced in 2010, the Netherlands' SDE+ programme was one of the first large multi-technology auctions of energy price support in Europe. All technologies except offshore wind compete within one auction, subject to reserve prices specific to each technology. In other cases, only comparable technologies compete. For example, the UK Contracts for Difference scheme divides technologies into groups (in the 4th round, these were "established", e.g., onshore wind and solar; "less established", e.g., advanced conversion technologies; and offshore wind), and auctions are run separately for each group.

[2] The mechanism that I consider (i.e., a Product-Mix Auction) automatically increases efficiency in my model in which each bidder demands just one of the goods, of which she demands a single unit, so the bidder therefore bids her true value. My model applies more generally if there are sufficiently many bidders that competitive bidding is a reasonable assumption in uniform-price auctions. This precludes the oligopolistic effects in multi-market models such as Chen and Duffie (2021), Rostek and Yoon (2021), Wittwer (2021), etc.

benefits while the other side's expected surplus falls.

My results are especially simple in the special case of my model in which the goods are symmetric. Both the bidders (in aggregate) and the auctioneer gain if the aggregate bid functions are neither extremely convex nor extremely concave. Otherwise, either the bidders or the auctioneer lose. For commonly used distributions, both bidders and the auctioneer gain from switching to the optimal mechanism. A notable exception is that bidders with constant-elasticity aggregate demand lose.

I also estimate the magnitude of these distributional effects in the primary market for Mexican government debt, using novel bid-level data from auctions held by the Mexican central bank to estimate the bidders' and auctioneer's preferences. I compare running the efficient multi-product auction for a pair of zero-coupon bonds (known as CETES, the acronym for Treasury Certificates) to running the existing pair of independent auctions, assuming participants bid competitively. I find the efficient auction of 182-day and 364-day bonds increases both bidder and auctioneer surplus, by 0.61 basis points of the issuance (or MEX$1.44 million) and 0.36 basis points (MEX$0.84m), respectively, over 2017–2021. The efficient auction of 28-day and 91-day bonds also increases both surpluses, by 0.08 basis points (MEX$0.13m) and 0.13 basis points (MEX$0.22m), respectively.

The closest work to mine is Fabra and Montero (2023). They analyse several mechanisms that allocate multiple units of different goods, assuming linear aggregate bid functions and linear marginal benefits for the auctioneer. They compare running a pair of independent auctions, in which the quantities of the two goods are fixed in advance to maximise expected welfare, with a multi-product auction in which the goods are treated interchangeably.[3] This chapter allows arbitrary aggregate bid functions and auctioneer costs of production. Also going beyond Fabra and Montero (2023), I obtain results comparing the distributional effects of the independent auctions to those of the optimal mechanism.[4]

Section 2 describes an auctioneer who wishes to allocate a fixed total quantity across two goods that she views as imperfect substitutes, for example two bonds with different maturities or two alternative energy sources. The objective is to maximise welfare, as conventionally defined by the sum of total bidder and auctioneer surplus.[5] For each good, there is a large

[3]Fabra and Montero (2023) define welfare by a weighted sum of the sum of bidder and auctioneer surplus, and total revenue. I focus on the case in which the weighting on revenue is zero.

[4]Fabra and Montero (2023) illustrate their model's predictions for Spain's 2021 auction of wind and solar energy, which I also analyse in Appendix H. I thank Natalia Fabra and Juan-Pablo Montero for kindly sharing their simulation estimates.

[5]Section 6.1 analyses the case in which the objective is to maximise auctioneer surplus.

number of risk-neutral bidders with unit demand, who face perfectly correlated shocks to their values.[6] It follows that the efficient auction is a Product-Mix Auction (PMA) as described in Klemperer (2008).[7] The PMA finds the competitive equilibrium, assuming bidders submit bids equal to their true values and that the auctioneer expresses her true costs of producing the different goods. In my model, it is indeed rational for all participants to behave "truthfully" in this way, so that the PMA maximises ex post welfare.[8]

The main contribution of this chapter is to compare the effects on bidder and auctioneer surpluses of this optimal mechanism to running a pair of separate simultaneous auctions (SSA). In the SSA, the quantity sold of each good is chosen in advance to maximise expected welfare, and so is insensitive to the shocks to bidders' values. All sales mechanisms are uniform-price.[9] The PMA increases total welfare by adjusting its allocation to the shocks, but the gain is not necessarily shared by both sides of the market.

Section 3 shows that we can decompose the difference in expected bidder surplus between the PMA, i.e., the "efficient auction", and the SSA into two effects. The Bidder Uncertainty Effect captures the impact of the randomness of the quantity sold of each good in the efficient auction. With this randomness, more bidders win when the price is lower. So bidders as a group gain if the randomness does not affect the average market-clearing price. This is just a reinterpretation of the classical result that a consumer with downward-sloping demand benefits from a mean-preserving spread in price (Waugh, 1944). Bidders *a fortiori* gain if the randomness reduces the average price. If the average price instead rises, bidders gain only if the benefits of the mean-preserving spread dominate. Specifically, bidders gain (lose) from the randomness if the aggregate bid function is log concave (log convex) (Proposition 1).

The second effect, the Bidder Allocation Effect, captures the impact on expected bidder surplus of changes in the expected allocations of the goods. It is positive (negative) if, from the perspective of the bidders, the efficient auction increases the expected quantity sold of the relatively more (less) profitable good (Proposition 2).

[6] With a large number of bidders, idiosyncratic shocks would be unimportant (and in the limit, irrelevant).

[7] I consider a simpler version of a PMA, in which bidders are not permitted to express their preferences between goods in the form of "paired bids", than Klemperer (2008, 2010, 2018). In other words, I consider a "one-sided" PMA. However, the two mechanisms are equivalent in my setting because the bidders view the two goods as non-substitutable.

[8] My model also applies if a large number of bidders for each good have identical downward-sloping demand curves, since it is rational for bidders who are small relative to the market to bid truthfully.

[9] In my competitive framework, discriminatory pricing would give identical results in equilibrium providing the bidders observe the shocks to both goods.

The impact of the efficient auction on the auctioneer is more complex than on the bidders for two reasons. First, the auctioneer's marginal costs of producing the two goods may not be independent so her surpluses from selling the two goods cannot be analysed separately. Second, the marginal winning bidder sets the auction price so, on the margin, bidder surplus is zero but auctioneer surplus may be different from zero.

Section 4 shows that it is therefore helpful to divide the difference in auctioneer surplus between the efficient auction and the SSA into two effects, which distinguish the impact on the price level from the impact on the relative prices. The first effect, the Auctioneer Price Level Effect, is positive (negative) if the efficient auction increases (decreases) the expected price level, as measured by the price of one of the goods. Specifically it is positive (negative) if both the aggregate bid function for that good is convex (concave) and the efficient auction decreases (increases) the expected quantity sold of that good (Proposition 3).

The second effect, the Auctioneer Relative Price Effect, captures the impact of the efficient auction on the relative prices. There are two aspects to this. The first corresponds to the Bidder Uncertainty Effect—the auctioneer gains from the uncertainty that the efficient auction generates in the relative prices. The second reflects the impact of the average difference in relative prices between the efficient auction and the SSA. Overall, the Auctioneer Relative Price Effect is positive (negative) if both the auctioneer's inverse marginal cost of one good relative to another is log concave (log convex) and the efficient auction increases (decreases) the good's expected quantity sold (Proposition 4).

Section 5 considers special cases of the model. In Section 5.1, I focus on the case in which both the bidders' aggregate preferences, and the auctioneer's preferences, are symmetric in the two goods. The distributional effects are particularly clean. If the aggregate bid functions are convex but log concave, the efficient auction increases the expected surplus of both the bidders and the auctioneer. Bidders as a group gain from the uncertainty in the quantities sold in the efficient auction, but this uncertainty also translates into higher expected auctioneer revenue. If the aggregate bid functions are either concave or very convex (i.e., log convex), only one side of the market benefits in expectation. Section 5.2 examines the asymmetric case of the generalised Pareto distribution, special cases of which correspond to the more commonly used functional forms of demand: linear, log linear, and constant elasticity.

Section 6 extends the model. When the objective is to maximise auctioneer surplus, the auctioneer of course benefits from the PMA and the conditions for the bidders to gain and to lose are similar to those in the main model (Section 6.1). Section 6.2 shows how

the distributional effects may differ with a variable total supply. Section 6.3 analyses an alternative multi-product design, commonly used in quantitative easing.[10] This "reference price auction" (RPA) imposes a fixed price difference between the goods. Because the optimal mechanism is between the extremes of fixing quantities, i.e., an SSA, and fixing relative prices, i.e., an RPA, it is not surprising that the results of the comparison of the efficient auction to the RPA are broadly opposite to that of the SSA.

In Section 7, I estimate the magnitudes of the distributional effects using bid-level data from the Mexican primary market for government debt. I group the four bonds that I study into two pairs: 28- and 91-day CETES, and 182- and 364-day CETES. For each pair, I find that the efficient auction increases both average bidder and average auctioneer surplus relative to the existing SSA design. The effects for each pair of bonds can be understood in terms of the model. For the 182- and 364-day pair, the aggregate bid functions are log concave, so the bidders gain from the uncertainty in issuance and the Bidder Uncertainty Effect is positive. This dominates the negative, albeit statistically insignificant, Bidder Allocation Effect. For the 28- and 91-day pair, the positive effect on bidder surplus is driven by a positive Bidder Allocation Effect, which results from the efficient auction increasing the expected allocation of the bond which is relatively more profitable on the margin for the bidders. The Bidder Uncertainty Effect is also positive, but statistically insignificant. For the auctioneer, the Auctioneer Relative Price Effect is zero under the assumptions of our model so the results are purely a function of the Auctioneer Price Level Effect. It is positive for the 28- and 91-day pair because the efficient auction increases the price of the reference good. It is also positive, albeit imprecisely estimated, for the 182- and 364-day bonds.

Section 8 concludes. Proofs of the results are found in Appendix A.

2 Model

An auctioneer allocates a single unit across two divisible goods, $j = \{1, 2\}$. Her total cost of producing q_1 of good 1 and q_2 of good 2 is $C(q_1, q_2)$, which is three-times continuously differentiable. It is convenient to assume her marginal cost of good $j = \{1, 2\}$, $MC_j(q_1, q_2) = \frac{\partial C(q_1,q_2)}{\partial q_j}$, is strictly increasing and that the goods are strictly imperfect substitutes, so that $\frac{\partial^2 C(q_1,q_2)}{\partial q_j \partial q_k} < \frac{\partial^2 C(q_1,q_2)}{\partial q_j^2}, k \neq j$. We do this for convenience to permit invertibility, but the cases in which the marginal cost curves have zero slopes or the goods are perfect substitutes follow automatically.

[10] Both the Federal Reserve and Bank of England used this design in their quantitative easing programmes (Bank of England, 2022; Song and Zhu, 2018).

The auctioneer always sells the entire unit.[11] It is therefore helpful to also define the auctioneer's *relative* marginal cost of producing good 1 (with reference to good 2), $\widetilde{MC}_1(q_1) = \frac{\partial C(q_1,1-q_1)}{\partial q_1} - \frac{\partial C(q_1,1-q_1)}{\partial q_2}$. This expresses the net marginal cost of increasing the quantity allocated of good 1 and correspondingly reducing the quantity of good 2, so that one unit is allocated in total. We can define the relative marginal cost of good 2 (with reference to good 1) similarly.

Since marginal costs are strictly increasing, her relative marginal cost is strictly increasing. We can therefore also define a corresponding inverse relative marginal cost curve, $\widetilde{mc}_1(\widetilde{p}_1)$, which specifies the quantity that the auctioneer is willing to sell of good 1 as a function of its relative price (equal to $\widetilde{p}_1 = p_1 - p_2$).[12] I assume $\lim_{\widetilde{p}_1 \to -\infty} \widetilde{mc}_1(\widetilde{p}_1) = 0$ and $\lim_{\widetilde{p}_1 \to \infty} \widetilde{mc}_1(\widetilde{p}_1) = 1$.

Good $j = \{1, 2\}$ is demanded by a measure, greater than one, of infinitesimal, risk-neutral bidders, each with value for an increment dq_j and quasi-linear utility. Their values are sums of idiosyncratic components, capturing preference heterogeneity among the bidders, and a correlated one, interpreted as a common taste shock to bidders for good j. Their idiosyncratic values are distributed by the twice continuously differentiable and strictly increasing function F_j with density function f_j.[13] The bidders for good j face a shock, θ_j, to their values, which is distributed on $[\underline{\theta}_j, \bar{\theta}_j]$ with zero mean. The vector of shocks is denoted by $\boldsymbol{\theta} = (\theta_1, \theta_2)$.

In auctions with uniform-pricing, a bidder with unit demand has a dominant strategy to submit a bid equal to his value. It follows that bidders in our setting also have this dominant strategy, which I refer to as 'bidding truthfully'. The aggregate bid function for good j — its demand curve — is therefore $d_j(p_j, \theta_j) = (1 - F_j(p_j - \theta_j))$, which is continuous and strictly decreasing in p_j and only depends on the difference between p_j and θ_j. We can also define a corresponding inverse demand curve, $D_j(q_j, \theta_j)$, which is additive in the shock θ_j. I describe $d_j(p_j, \theta_j)$ as log convex (log concave) if it is log convex (log concave) with respect to p_j, holding θ_j constant, i.e., $\frac{\partial^2 log(d_j(p_j,\theta_j))}{\partial p_j^2} \geq (\leq)\ 0$. Similarly, I describe $D_j(q_j, \theta_j)$ as convex (concave) if $\frac{\partial^2 D_j(q_j,\theta_j)}{\partial q_j^2} \geq (\leq)\ 0$.

Total bidder surplus, denoted $TBS(q_1, \boldsymbol{\theta})$, equals the total gross value to bidders of receiving

[11] I assume the submitted bids are always sufficiently high relative to the costs of production that the auctioneer never prefers to sell less than the full unit. Section 6.2 allows the total quantity sold to vary.

[12] The inverse relative marginal cost curve for good 2 is similarly defined and can be equivalently expressed by $\widetilde{mc}_2(\widetilde{p}_2) = (1 - \widetilde{mc}_1(-\widetilde{p}_2))$.

[13] With a large number of bidders, idiosyncratic shocks would be unimportant and, in the limit, irrelevant to the analysis: the outcome would be identical and, because bidders are risk-neutral, they collectively value the auction outcome as would a single representative consumer, so the measurement of bidder surplus would be unchanged.

q_1 of good 1 (and $(1-q_1)$ of good 2), less the total payment to the auctioneer, given the shocks $\boldsymbol{\theta}$. Analogously, total auctioneer surplus, $TAS(q_1, \boldsymbol{\theta})$, equals the total payment received by the auctioneer, minus her total cost of producing the two goods.

The auctioneer can adopt either of two sales mechanisms: a Product-Mix Auction (PMA) or a pair of separate simultaneous auctions (SSA).[14] In this context, the PMA is a multi-product auction in which bidders submit bids for the two goods and the auctioneer expresses her preferences between the goods.[15] It uses the information from the resulting demand and supply curves to find the competitive equilibrium, under the assumption that all bidders bid truthfully and the auctioneer expresses her true costs of producing the goods. Truthful behaviour is indeed rational for all participants as winning bidders of a good pay the highest losing bid for that good. The PMA therefore maximises ex post welfare, defined conventionally by the sum of total bidder surplus and auctioneer surplus. I refer to the PMA as the "efficient auction".

In the SSA, the quantity sold of each good is fixed ex ante and chosen by the auctioneer to maximise expected welfare. Bidders submit the same bids (for each good) as they submit in the efficient auction, because it is optimal to bid truthfully and bidders view the two goods as non-substitutable. The two auction prices equal the highest losing bids on the corresponding goods.

The quantity sold of good 1 under mechanism $m \in \{S, E\}$, where $m = S$ refers to the SSA and $m = E$ to the efficient auction, is q_1^m (and $q_2^m = 1-q_1^m$), where q_1^S is deterministic and q_1^E is a function of $\boldsymbol{\theta}$, and $p_j(q_j^m, \theta_j) = D_j(q_j^m, \theta_j)$ is good j's auction price. The efficient auction finds the quantities such that the relative value of the goods to the marginal winning bidders equals the goods' relative marginal cost, i.e., $D_1(q_1^E(\boldsymbol{\theta}), \theta_1) - D_2(1-q_1^E(\boldsymbol{\theta}), \theta_2) = \widetilde{MC}_1(q_1^E(\boldsymbol{\theta}))$. The fixed quantities in the SSA satisfy this in expectation, $\mathbb{E}\left[D_1(q_1^S, \theta_1) - D_2(1-q_1^S, \theta_2)\right] = \widetilde{MC}_1(q_1^S)$.[16]

[14] I consider an alternative design, in which the price difference between goods is fixed in advance, in Section 6.3.

[15] This is a "one-sided" PMA because only the auctioneer expresses preferences between goods. In the general PMA, both the bidders and the auctioneer are permitted to do so, with bidders submitting "paired bids" (see Klemperer, 2008, 2010, 2018). Because bidders view the two goods as non-substitutable, the one-sided PMA and general PMA are equivalent in my setting.

[16] The SSA allocation maximises expected welfare and so solves

$$\max_{q_1^S} \mathbb{E}\left[\int_0^{q_1^S} D_1(q_1, \theta_1) dq_1 + \int_0^{1-q_1^S} D_2(q_2, \theta_2) dq_2 - \int_0^{q_1^S} \widetilde{MC}_1(q_1) dq_1\right]$$

The relative price of good 1 is $\widetilde{p}_1(q_1^m, \boldsymbol{\theta}) = p_1(q_1^m, \theta_1) - p_2(1 - q_1^m, \theta_2)$. Let $\underline{q}_1 = q_1^E((\underline{\theta}_1, \bar{\theta}_2))$ and $\bar{q}_1 = q_1^E((\bar{\theta}_1, \underline{\theta}_2))$ denote the minimum and maximum quantities sold of good 1 in the efficient auction, and $(\underline{q}_2, \bar{q}_2) = (1 - \bar{q}_1, 1 - \underline{q}_1)$. I assume that these are interior solutions.

I label the goods so that the efficient auction increases the expected quantity sold of good 1, that is, $\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right] \geq q_1^S$. A sufficient condition for this labelling is that inverse demand for good 1 relative to good 2 is weakly more convex than the relative marginal cost of good 1 (see Appendix B).

3 Bidder surplus

The expected surplus of a winning bidder of good $j = \{1, 2\}$ is the difference between his expected value and the good's expected auction price, conditional on him winning. His probability of winning and the conditional expected price differ between the two sales mechanism and this difference depends on the distribution of shocks and the bidder's own value. In general, some bidders gain and some lose from the efficient auction relative to the SSA,[17] so I sum the effects across bidders.[18]

Measuring (total) bidder surplus as a function of quantities, rather than prices as is conventional, is helpful.[19] Bidder surplus for good $j \in \{1, 2\}$ is the sum of the values of winning bidders, less their payment to the auctioneer. The difference in the value of the winners is simply the area under the inverse demand curve between the quantities allocated in the efficient auction and SSA. That is, it equals the sum of Areas L_j, M_j, and N_j in Figure 1 for shocks $\boldsymbol{\theta}$.

Similarly, the difference in auctioneer revenue is measured by the area under the marginal revenue curve for good j, $MR_j(q_j, \theta_j) = \frac{\partial(D_j(q_j, \theta_j)q_j)}{\partial q_j}$, because this curve describes the change in total payment for good j due to a marginal increase in its quantity sold. This is equal to Area N_j and is equivalent to the difference in the product of price and quantity sold, i.e., (Area M_j + Area N_j) − (Area J_j + Area K_j).

The difference in total bidder surplus for good $j \in \{1, 2\}$ is therefore the area between the inverse demand curve and the marginal revenue curve between the quantities sold in the

[17] For example, a bidder with a low value always loses in the SSA but wins with positive probability in the efficient auction, and so unambiguously gains. In contrast, a bidder with a high value always wins in the SSA but sometimes loses in the efficient auction. The efficient auction also changes the expected price he pays, conditional on winning. If the demand curve is concave, he gains. If the demand curve is convex, the net effect on his expected surplus depends on the relative importance of these two opposing effects.

[18] This is of course equivalent to the average effect on expected individual bidder surplus.

[19] Bulow and Klemperer (2012) demonstrate the value of measuring surplus in this way.

efficient auction and SSA. For shocks $\boldsymbol{\theta}$, it is the sum of Areas L_j and M_j.

Another way to understand this is that the marginal winning bidder is indifferent to winning because his bid sets the auction price. So bidder surplus on the margin is simply the change in total payment of the inframarginal winners, i.e., $D_j(q_j, \theta_j) - MR_j(q_j, \theta_j) = -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j$. For shocks $\boldsymbol{\theta}$, the difference in total bidder surplus for good $j \in \{1, 2\}$ is therefore

$$\int_{q_j^S}^{q_j^E(\boldsymbol{\theta})} (D_j(q_j, \theta_j) - MR_j(q_j, \theta_j)) \ dq_j = \int_{q_j^S}^{q_j^E(\boldsymbol{\theta})} -\frac{\partial D_j(q_j, \theta_j)}{\partial q_j} q_j \ dq_j.$$

I decompose the difference in expected total bidder surplus between the efficient auction and SSA into two effects. Let

$$\text{Bidder Uncertainty Effect} = \ \mathbb{E}\left[\sum_{j=\{1,2\}} \left(\int_{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}^{q_j^E(\boldsymbol{\theta})} -\frac{\partial D_j(q_j, \theta_j)}{\partial q_j} q_j \ dq_j \right) \right]$$

and

$$\text{Bidder Allocation Effect} = \ \mathbb{E}\left[\sum_{j=\{1,2\}} \left(\int_{q_j^S}^{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]} -\frac{\partial D_j(q_j, \theta_j)}{\partial q_j} q_j \ dq_j \right) \right]$$

The following lemma immediately follows.

Lemma 1 *The difference in expected total bidder surplus between the efficient auction and SSA,* $\mathbb{E}\left[TBS(q_1^E(\boldsymbol{\theta}), \boldsymbol{\theta}) - TBS(q_1^S, \boldsymbol{\theta})\right]$*, equals the sum of the Bidder Uncertainty Effect and the Bidder Allocation Effect.*

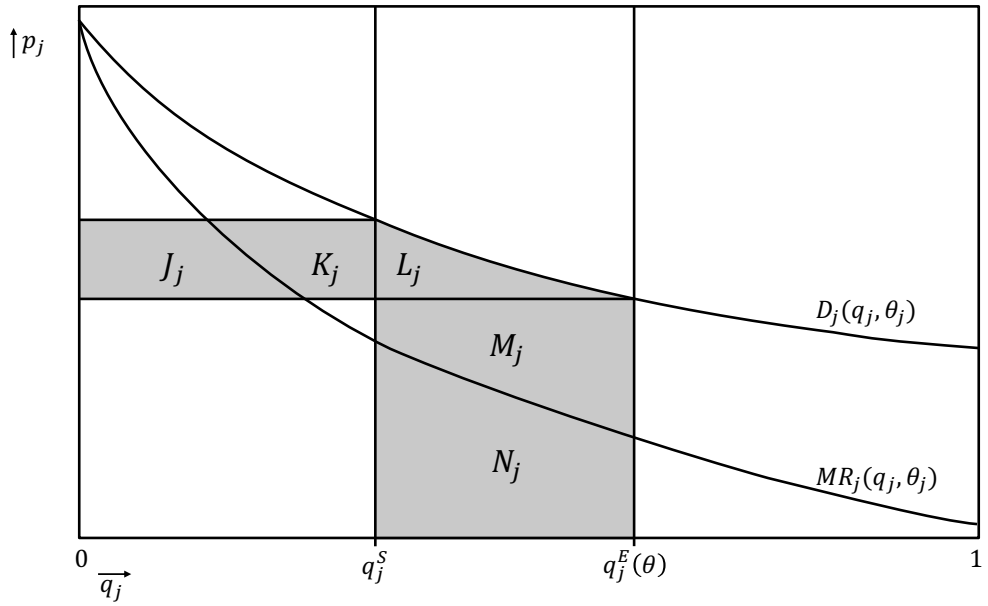
These effects each relate to properties of the bidders' demand curves and the auctioneer's costs of production.

3.1 Bidder Uncertainty Effect

The Bidder Uncertainty Effect captures the benefit to bidders of uncertainty in the quantities sold of the two goods in the efficient auction.

A marginal increase in the quantity sold of good $j = \{1, 2\}$ reduces the auction price for all winning bidders, and causes an additional bidder to win when he would otherwise lose. Only

Figure 1: Change in total bidder surplus for good $j \in \{1, 2\}$: two equivalent measurements



The difference in total bidder surplus between selling $q_j^E(\boldsymbol{\theta})$ and q_j^S is measured by the sum of Areas L_j and M_j. This is equivalent to the conventional measure of the change in bidder surplus, which is the sum of Areas J_j, K_j, and L_j. This is evident by comparing two equivalent measures of total payment. The conventional measure of the change in total payment is (Area M_j + Area N_j) **less** (Area J_j + Area K_j). This is equivalent to the area under the marginal revenue curve, i.e., Area N_j. Hence, Area M_j = (Area J_j + Area K_j).

the first effect is relevant to bidder surplus, as the marginal winner is indifferent between winning and losing, paying a price precisely equal to his value $D_j(q_j, \theta_j)$.

In the simple case of linear demand, bidders in aggregate benefit from uncertainty, just as a consumer with downward-sloping demand benefits from a mean-preserving spread in price in Waugh's (1944) classical result. This is because equally sized changes in the quantity sold of a good cause equally sized changes in its auction price, and more bidders win when the auction price is low. With non-linear demand, the changes in auction price are asymmetric. Uncertainty unambiguously benefits the bidders if the demand curve is concave: when the quantity sold increases, the larger number of winners benefit from a greater reduction in price. But if the demand curve is sufficiently convex, the relatively steep price rise for the smaller number of winners, when the quantity sold declines, dominates the price fall from the increase in quantity. So bidders as a group lose out from the uncertainty.

Specifically, the Bidder Uncertainty Effect is positive if bidder surplus for good $j = \{1, 2\}$, $\int_0^q -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j \; dq_j$, is a convex function of the quantity sold, q. In this case, expected bidder surplus for good j in the efficient auction, i.e., $\mathbb{E}\left[\int_0^{q_j^E(\theta)} -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j \; dq_j\right]$ is greater than bidder surplus evaluated at the expected quantities sold, i.e., $\int_0^{\mathbb{E}\left[q_j^E(\theta)\right]} -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j \; dq_j$ (by Jensen's inequality), so it is immediate from its definition that the Bidder Uncertainty Effect is positive. Simple manipulation of $\int_0^q -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j \; dq_j$ shows that it is a convex function of q if and only if the demand curve is log concave.[20] So:

Proposition 1 *The Bidder Uncertainty Effect is positive (negative) if both demand curves are log concave (log convex).*[21]

A large number of well-known distributions have a log concave density function, f_j, which implies a log concave demand curve, $d_j(p_j, \theta_j)$.[22] Examples include uniform (corresponding to linear demand), normal, logistic, extreme value, and exponential (corresponding to log

[20] This can also be understood in terms of the fatness of the right tail of the distribution of bidders' values. With a fat tail (corresponding to a log convex demand curve), few bidders value the good extremely highly. If only a small quantity is sold, it will be sold only to these bidders and the auction price—corresponding to the marginal winner's value—will be very high. As more bidders win when larger quantities are sold, the price will decline only gradually as the values of the marginal winners become more similar.

[21] Naturally, it is sufficient that the condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j, \theta_j) : d_j(p_j, \theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j = \{1, 2\}$. The effect is strictly positive (negative) if at least one demand curve is strictly log concave (log convex).

[22] The survival function (in this context, the demand curve) of a continuously differentiable log concave density function is also log concave, but the converse does not necessarily hold for log convex densities (Bagnoli and Bernstrom, 2005).

linear demand).[23] Indeed, there are many arguments for why log concave demand might be a natural assumption. For example, Weyl and Fabinger (2013) argue that, if bidders' willingness to pay for the good were proportional to their income, then log concave demand fits the income distributions commonly observed in many countries.[24] A notable exception to the set of distributions with log concave density is Pareto (corresponding to constant elasticity demand).

There are two additional ways to understand this result. First, my model also applies to a situation in which all bidders for good j have identical downward-sloping demand curves, as described by the demand curve, $d_j(p_j, \theta_j)$, because it is also rational for bidders in this case to behave truthfully when there are sufficiently many of them. In this view, each bidder's surplus is measured by the area between the inverse demand curve and the marginal revenue curve. If a bidder is risk-loving (risk-averse) with respect to quantity, that is, he prefers (dislikes) uncertainty in the quantity sold, his demand curve is log concave (log convex).[25]

Second, the randomness in the quantities sold in the efficient auction is akin to the rationing created by price controls in Bulow and Klemperer (2012). In my context, for any given quantity sold, the price of each good reacts one-for-one to the common shock to bidders' values for that good, because it is equal to the marginal winner's value for that good. And so the bidder surplus does not depend directly on the shocks, and depends only on the winners' idiosyncratic values (i.e., values less the shock). The efficient auction creates uncertainty in the quantities sold, so it is in effect rationing bidders with medium idiosyncratic values. The reason is that the lower part of these medium-value bidders would always lose if the quantities sold were the expected quantities sold in the efficient auction but do sometimes win in the efficient auction, while the higher part of the medium-value bidders would always win if the quantities sold were the expected quantities but do sometimes lose. Bulow and Klemperer (2012) show that if the demand curves are log concave (log convex), bidders gain (lose) in aggregate from rationing if there is no reduction in supply. So it follows that, in our context, bidders gain (lose) from the uncertainty in the quantities sold generated by the

[23] See Bagnoli and Bernstrom (2005) for an extensive list.

[24] On the other hand, Fabinger and Weyl (2012) suggest that log convex demand might be a reasonable characterisation if bidders have heterogeneous valuations for a good but resale is possible. The ability to resell the good places a floor on bidders' willingness to pay and heterogeneous valuations mean that the demand curve is downward-sloping above this floor.

[25] The coefficient of relative risk aversion for an agent with random wealth w and expected utility $\mathbb{E}[u(w)]$ is $\frac{-wu''(w)}{u'(w)}$. If the representative bidder's expected utility is the area between the demand curve and the marginal revenue curve, then the coefficient of relative risk aversion for the bidder with random allocation q_j is $r(q_j) = \frac{-q_j}{(D_j(q_j,\theta_j)-MR_j(q_j,\theta_j))} \frac{\partial(D_j(q_j,\theta_j)-MR_j(q_j,\theta_j))}{\partial q_j}$. If the demand curve is log concave (log convex), $r(q_j) \leq 0$ ($r(q_j) \geq 0$) (see proof of Proposition 1).

efficient auction.

3.2 Bidder Allocation Effect

The Bidder Allocation Effect represents the impact on expected total bidder surplus of the difference between sales mechanisms in the expected quantity sold of each good. The efficient auction increases the expected quantity sold of good 1, so the Bidder Allocation Effect is positive (negative) if good 1 is relatively more (less) profitable on the margin for bidders.

The marginal winning bidder for each good is indifferent between winning and losing, so that only the change in surplus for inframarginal winners is relevant. Good 1 is therefore relatively more profitable if an increase in its quantity sold (and corresponding reduction in the quantity sold of good 2) reduces the average price paid by the inframarginal winners across the two goods.

Sufficient conditions for this are that good 1's expected auction price is higher than good 2's, i.e., if good 1 is relatively costly to produce, and demand for good 1 is less price elastic.[26] The reason is that a marginal increase in the quantity sold of good 1 causes a larger quantity-weighted change in price in this case.

Proposition 2 *The Bidder Allocation Effect is positive (negative) if both demand curves are log concave (log convex),*[27] *good 1's expected price elasticity of demand is smaller (larger) in absolute value than good 2's and the relative marginal cost of good 1 is positive (negative) at the quantities sold in the SSA.*[28]

Many frequently used demand curves are log concave (see Section 3.1). Moreover, if the demand curves are log concave, then if one of the other two conditions required for the Bidder Allocation Effect to be positive holds, it is more likely that the other holds. This is because if each demand curve is log concave, the absolute value of its price elasticity is decreasing in its quantity. And so, if good 1 is less price elastic than good 2 at the SSA allocation, the quantity sold of good 1 in the SSA is typically relatively large. In this case, the relative marginal cost of good 1 is typically positive, because it is increasing in the quantity

[26] It is of course also sufficient that both a larger quantity is sold of good 1 than of good 2 and the inverse demand curve for good 1 is steeper.

[27] As for Proposition 1, it is sufficient that this condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j, \theta_j) : d_j(p_j, \theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j = \{1, 2\}$.

[28] The inequality is strict if the efficient auction strictly increases the expected quantity sold of good 1 and one of the three stated conditions holds strictly.

sold of good 1. On the other hand, it seems less likely that the two conditions required for the Bidder Allocation Effect to be negative when demand is log convex would both hold.

Propositions 1 and 2 provide conditions under which we can determine the effect of switching to the efficient auction on expected bidder surplus.

Corollary 1 *The efficient auction increases (decreases) expected bidder surplus relative to the SSA if both demand curves are log concave (log convex), good 1's expected price elasticity of demand is smaller (larger) in absolute value than good 2's, and the relative marginal cost of good 1 is positive (negative) at the quantities sold in the SSA.*

3.3 Market-specific bidder surplus

The impact of the efficient auction on the expected surplus of bidders for good $j \in \{1, 2\}$ can be decomposed into good-specific Bidder Uncertainty and Allocation Effects. By the labelling of the two goods, the good-specific Bidder Allocation Effect is positive (negative) for good 1 (good 2), and the conditions for the good-specific Bidder Uncertainty Effect are a natural corollary to Proposition 1 (see Appendix C).

4 Auctioneer surplus

Auctioneer surplus is the difference between the total revenue raised in the auction and the auctioneer's total cost of producing the goods. The efficient auction increases expected surplus for at least one side of the market. If it reduces expected bidder surplus—as determined by the conditions in Section 3—it therefore increases expected auctioneer surplus. In this section, I provide additional conditions which determine the efficient auction's impact on expected auctioneer surplus. (Analogously, the efficient auction must increase expected bidder surplus if it reduces expected auctioneer surplus.)

Measuring the impact of the efficient auction on the auctioneer is more complicated than for the bidders for two reasons.

First, the auctioneer's marginal cost of each good may be strictly below its auction price because the total supply is fixed. This would create strictly positive surplus on the margin for the auctioneer, whereas the marginal winning bidder is always indifferent between winning and losing, paying an auction price equal to his own bid.

This is relevant even when the auctioneer has independent marginal costs of production, in which case the difference in auctioneer surplus between the efficient auction and SSA can be

measured conventionally by considering the two goods separately. This is illustrated for a realisation of shocks for which $q_1^E(\boldsymbol{\theta}) > q_1^S$ (and correspondingly $q_2^E(\boldsymbol{\theta}) < q_2^S$), in Figure 2 by (Area Q − Area P + Area R − Area S).

Second, with more general cost functions, the two goods cannot be analysed separately, because the marginal cost of producing one good is dependent on the quantity produced of the other.[29]

To address both of these issues, it is therefore helpful to decompose auctioneer surplus differently from bidder surplus in a way which isolates changes in the price level from changes in the relative price structure of the goods. The auctioneer's total revenue, $p_1(q_1, \theta_1)q_1 + p_2(1-q_1, \theta_2)(1-q_1)$, consists of the revenue that would be generated if she sold the entire unit of a "reference good" at its auction price (and none of the other good) and the additional revenue resulting from a part of the unit being reallocated to the other good. Measured with reference to good 2, the "reference revenue" is $p_2(1 - q_1, \theta_2)$ and the "relative revenue" is $\widetilde{p}_1(q_1, \boldsymbol{\theta})q_1$.

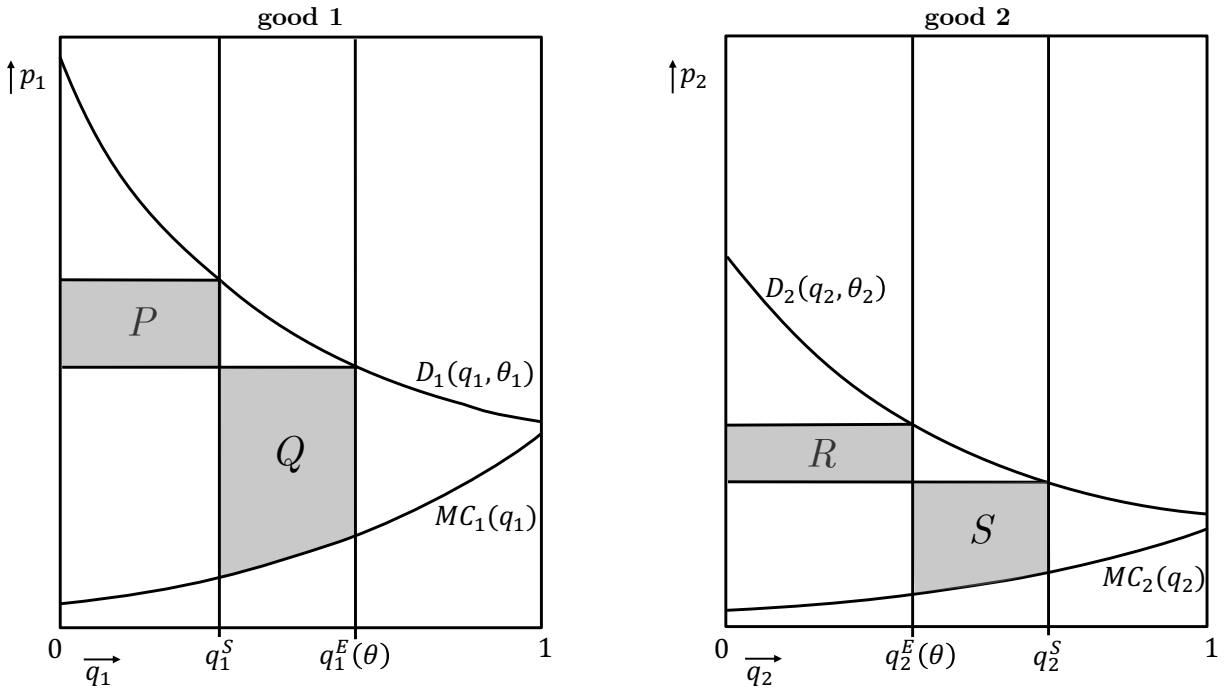
We can correspondingly decompose total cost, which is measured by the sum of the auctioneer's "reference cost" of producing the entire unit of good 2, $C(0, 1)$, and her "relative cost" of good 1, the area under her relative marginal cost curve, $\int_0^{q_1} \widetilde{MC}_1(q)dq$.

First, the auctioneer's "reference surplus"—the difference between her reference revenue and reference cost—is measured by $(p_2(1 - q_1, \theta_2) - C(0, 1))$. Given shocks $\boldsymbol{\theta}$ for which $q_1^E(\boldsymbol{\theta}) > q_1^S$, the difference in reference surplus between the efficient auction and SSA is measured by Area U in Figure 3.

Second, as for bidder surplus, the auctioneer's "relative surplus"—the difference between her relative revenue and relative cost—can be expressed as a function of quantities. The auctioneer's relative cost is measured by the area under the relative marginal cost curve and relative revenue is measured accordingly. That is, the auctioneer's relative marginal revenue curve is given by $\widetilde{MR}_1(q_1) = \frac{\partial(\widetilde{MC}_1(q_1)q_1)}{\partial q_1}$, which is the change in relative revenue due to a marginal increase in the quantity of good 1 (and corresponding reduction in good 2) if the relative price is set equal to the relative marginal cost. The auctioneer's relative revenue is the area under this curve, up to the quantity allocated of q_1, plus an adjustment, $(\widetilde{p}_1(q_1, \boldsymbol{\theta}) - \widetilde{MC}_1(q_1))q_1$, which reflects the difference between the relative price and the

[29] In Section 6.2, I show that auctioneer surplus can be analysed identically to bidder surplus in the case that there are both independent marginal costs (so that the two goods can be analysed separately) and variable total supply (so that the auctioneer's marginal cost for each good equals its auction price).

Figure 2: Conventional measure of auctioneer surplus with independent marginal costs



If the auctioneer's costs of producing the two goods are independent, the difference in auctioneer surplus between selling $q_1^E(\boldsymbol{\theta})$ units of good 1 (and corespondingly $q_2^E(\boldsymbol{\theta})$ units of good 2) and selling q_1^S units of good 1 (and correspondingly q_2^S units of good 2), given the realisation of shocks $\boldsymbol{\theta}$, is (Area Q − Area P + Area R − Area S).

auctioneer's relative marginal cost.

And so her relative surplus is equal to

$$\widetilde{p}_1(q_1, \boldsymbol{\theta})q_1 - \int_0^{q_1} \widetilde{MC}_1(q)dq = \int_0^{q_1} \left(\widetilde{MR}_1(q) - \widetilde{MC}_1(q)\right) dq_1 + \left[\widetilde{p}_1(q_1, \boldsymbol{\theta}) - \widetilde{MC}_1(q_1)\right] q_1$$

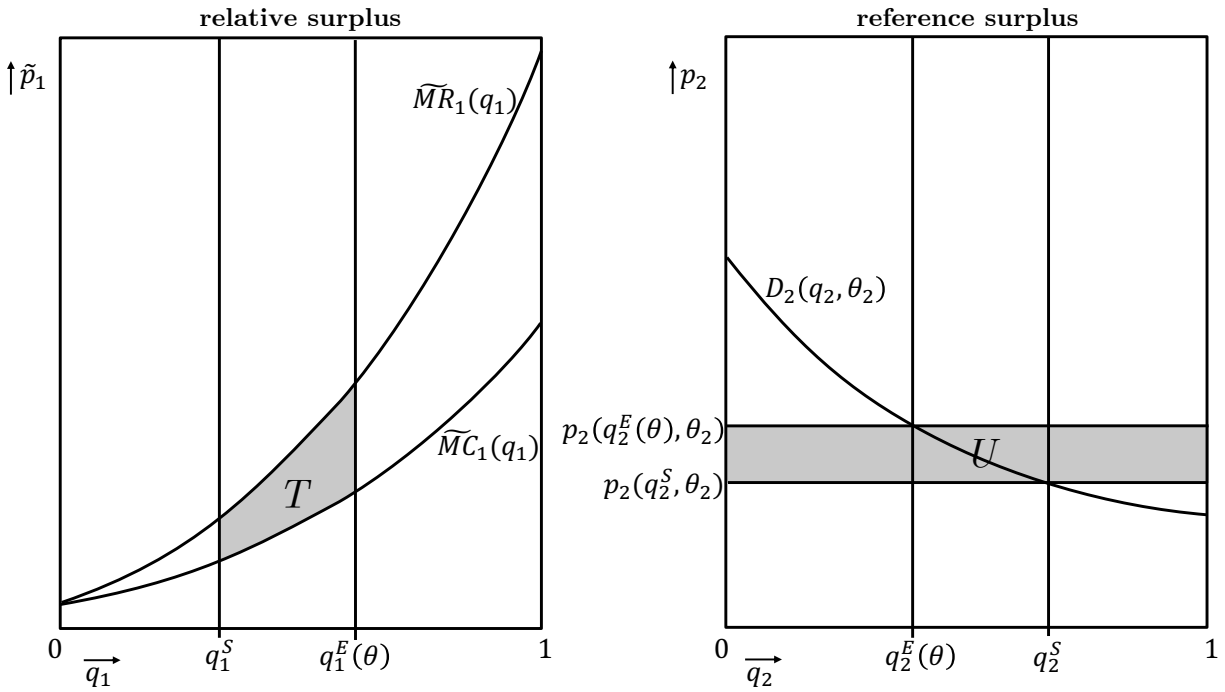
Given shocks $\boldsymbol{\theta}$ for which $q_1^E(\boldsymbol{\theta}) > q_1^S$, the difference in relative surplus between the efficient auction and SSA is measured by the sum of Area T in Figure 3 and the adjustment $(\theta_2 - \theta_1)q_1^S$.

The adjustment is zero in the efficient auction and zero in expectation in the SSA, and so is irrelevant to the difference in expected surplus between sales mechanisms.

Equivalently, we can decompose total revenue and total cost with reference to good 1 to define the reference and relative surpluses accordingly.[30]

[30]See Appendix D for the relationship between the areas in Figures 2 and 3, and for an illustration of the relative surplus and reference surplus with reference to good 1.

Figure 3: Auctioneer surplus measured with reference to good 2



The Auctioneer Relative Price Effect with reference to good 2 is the expected difference in relative surplus between the efficient auction and the SSA, that is, the expected size of Area T. The Auctioneer Price Level Effect with reference to good 2 is the expected difference in reference surplus between the efficient auction and the SSA, that is, the expected size of Area U.

I use these measures to decompose the difference in expected total auctioneer surplus between the efficient auction and SSA into two effects, the Auctioneer Price Level Effect and the Auctioneer Relative Price Effect, which respectively measure the auctioneer's reference surplus and relative surplus with reference to good $j = \{1, 2\}$:

$$\text{(Auctioneer Price Level Effect)}_j = \mathbb{E}\left[p_j(q_j^E(\boldsymbol{\theta}), \theta_j) - p_j(q_j^S, \theta_j)\right]$$
$$\text{and (Auctioneer Relative Price Effect)}_j = \mathbb{E}\left[\int_{q_k^S}^{q_k^E(\boldsymbol{\theta})} \left(\widetilde{MR}_k(q_k) - \widetilde{MC}_k(q_k)\right) dq_k\right]$$

where $k \neq j, k \in \{1, 2\}$.

Lemma 2 *The difference in expected auctioneer surplus between the efficient auction and SSA, $\mathbb{E}\left[TAS(q_1^E(\boldsymbol{\theta}), \boldsymbol{\theta}) - TAS(q_1^S, \boldsymbol{\theta})\right]$, equals the sum of the Auctioneer Price Level Effect and the Auctioneer Relative Price Effect, where the decomposition is made with reference to either good.*

These effects each relate to properties of the bidders' demand curves and the auctioneer's costs of production.

4.1 Auctioneer Price Level Effect

The Auctioneer Price Level Effect captures the impact of the efficient auction on expected auctioneer surplus through its impact on the price level, measured by the auction price of the reference good. Both the uncertainty in the quantities sold in the efficient auction and the difference in the expected quantities sold between mechanisms, are relevant.

If the demand curve of the reference good is convex (concave), uncertainty in the quantities sold increases (reduces) its expected auction price, increasing (reducing) expected revenue for the auctioneer and therefore expected auctioneer surplus.

Moreover, by the labelling of the two goods, the expected quantity sold of good 2 (good 1) is smaller (larger) in the efficient auction than the SSA, increasing (reducing) the reference good's expected auction price, i.e., the price level, when measured with reference to good 2 (good 1).

Proposition 3 *The Auctioneer Price Level Effect with reference to good 2 (good 1) is positive (negative) if the demand for good 2 (good 1) is convex (concave).*[31]

[31] Naturally, it is sufficient that the condition holds only for the range of feasible prices and shocks, that

4.2 Auctioneer Relative Price Effect

The Auctioneer Relative Price Effect captures the impact of the variation in the relative price structure caused by the efficient auction on the auctioneer's expected surplus.[32]

The efficient auction chooses an allocation at which, given the shocks, the relative price of good 1 equals its relative marginal cost. In the same way that the marginal winning bidder is indifferent between winning and losing, the auctioneer is indifferent on the margin between additional units of the two goods. Moreover, the SSA maximises expected welfare by setting the expected relative price equal to the relative marginal cost. The deviations of the relative price from relative marginal cost in the SSA therefore net out in expectation – the impact of the efficient auction can be measured as if the relative prices in the SSA and the efficient auction were set equal to the relative marginal cost of the quantities sold in the SSA and the relative marginal cost of the quantities sold in the efficient auction, respectively.

Because the efficient auction sets a relative price for good 1 equal to its relative marginal cost, which is an increasing function, more of good 1 is sold in the efficient auction when its relative price is higher. There are two implications of this.

First, uncertainty in the relative price has an effect on the auctioneer which is analogous to the Bidder Uncertainty Effect, with the auctioneer's inverse relative marginal cost curve, $\widetilde{mc}_1(\widetilde{p}_1)$, corresponding to the bidders' demand curve, $d_1(p_1)$. That is, the auctioneer benefits from a mean-preserving spread in the relative price, and therefore also benefits if the randomness *increases* the average relative price, i.e., if the inverse relative marginal cost curve is concave. It is only if the inverse relative marginal cost curve is sufficiently convex (specifically, if it is log convex) that the uncertainty causes the average relative price to fall so far that the auctioneer loses from the uncertainty in the relative price.

The second implication is that *ceteris paribus* the auctioneer prefers a higher average relative price. So the larger (smaller) expected quantity sold of good 1 (good 2) in the efficient auction relative to the SSA is associated with a higher (lower) relative marginal cost of good 1 (good 2), increasing (reducing) the auctioneer's relative surplus with reference to good 1 (good 2).

---

is $\forall\{(p_j, \theta_j) : d_j(p_j, \theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j \in \{1, 2\}$.

The inequality is strict if the efficient auction strictly increases the expected quantity sold of good 1 relative to the SSA and the condition on the convexity of demand holds strictly.

[32] The relative prices of the goods also vary in the SSA, as the price of each good reacts one-for-one to the shock to demand for that good. Because the quantities sold in the SSA are fixed, this variation has no effect on expected auctioneer surplus.

Proposition 4 *The Auctioneer Relative Price Effect with reference to good 2 (good 1) is positive (negative) if the inverse relative marginal cost curve of good 1 (good 2) is log concave (log convex).*[33]

A simple example of a positive Auctioneer Relative Price Effect with reference to good 2 is when the auctioneer's marginal costs are independent, and the slopes of the two inverse marginal cost curves are log concave.[34] As discussed in Section 3.1, many commonly used distributions have log concave densities, including uniform, normal, logistic, extreme value, and exponential, which correspond to log concave slopes if the inverse marginal cost curves are viewed as distribution functions.

Propositions 3 and 4 provide conditions under which we can determine the effect of switching to the efficient auction on expected auctioneer surplus.

Corollary 2 *The efficient auction increases (decreases) expected auctioneer surplus relative to the SSA if the inverse relative marginal cost curve of good 1 (good 2) is log concave (log convex) and the demand for good 2 (good 1) is convex (concave).*

5 Special cases

I first show that the results are particularly clean in the case in which the two goods are symmetric: the impact of the efficient auction on expected bidder and auctioneer surpluses depends on whether the demand curves are either concave (and therefore log concave), log concave but convex, or log convex (and therefore convex). I then show that these results also largely extend to asymmetric goods when the demand curves have standard functional forms. Finally, I discuss the case in which the auctioneer's marginal cost of one good relative to another is constant. This case will be important for our empirical analysis in Section 7.

[33]The effect is strictly positive (negative) if the relative marginal cost curve is strictly log concave (log convex).

It is of course sufficient that the respective conditions hold only for the range of feasible quantities sold, that is $\forall q_1 \in [\underline{q}_1, \bar{q}_1]$ where $q_2 = 1 - q_1$.

[34]With this property, the auctioneer's supply of each good becomes more elastic as the quantity increases. So uncertainty causes the relative price of good 1 to rise disproportionately at larger quantities sold of good 1. Moreover, the larger expected quantity sold of good 1 in the efficient auction increases the average relative price of good 1. So both implications of the fact that the relative price equals the relative marginal cost of good 1 in the efficient auction are positive. See Appendix E for details.

5.1 The symmetric case

Suppose the expected inverse demand curves are identical, i.e., $\mathbb{E}\left[D_1(q,\theta_1)\right] = \mathbb{E}\left[D_2(q,\theta_2)\right]$ $\forall q = [0,1]$; the shocks are exchangeable, i.e., the distribution of (θ_1,θ_2) is the same as the distribution of (θ_2,θ_1);[35] and the auctioneer's total cost of production is symmetric in the goods, i.e., $C(x,y) = C(y,x)$ $\forall x,y \in [0,1]$.

This model, which I refer to as the symmetric case, is instructive. We can split the space of possible demand curves into those for which the bidders and auctioneer share the welfare gain of the efficient auction (that is, both expected bidder and auctioneer surplus increase), and the case in which only one side of the market benefits in expectation.

The Bidder Allocation Effect is zero because the expected quantities sold of the two goods are equal both in the SSA and in the efficient auction in the symmetric case (i.e., $q_1^S = \mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right] = \frac{1}{2}$). So the impact of the efficient auction on expected bidder surplus only depends on the benefit to the bidders of uncertainty. It follows immediately from Proposition 1 that if the demand curves are log concave (log convex), the Bidder Uncertainty Effect is positive (negative) so that the efficient auction increases (decreases) expected bidder surplus.

The effects for the auctioneer are also particularly clean because we can show that the auctioneer unambiguously benefits from the impact of the efficient auction on the relative prices. Positive and negative deviations of the quantity sold of an arbitrary good from its expected value impact the good's relative price symmetrically. The efficient auction chooses an allocation for which the relative price is equal to the auctioneer's relative marginal cost, given the shocks. Since relative marginal cost is an increasing function, the auctioneer sells more of the good when its relative price is higher and *ceteris paribus* prefers selling the good at a higher relative price. So the Auctioneer Relative Price Effect is positive.[36]

The sign of the Auctioneer Price Level Effect is determined by the conditions of Proposition 3. If the demand curves are convex, the efficient auction therefore unambiguously increases expected auctioneer surplus. But if the demand curves are concave, the auctioneer loses from the efficient auction's impact on the price level but gains from its effect on relative prices. The efficient auction therefore increases expected auctioneer surplus relative to the SSA if the demand curves are either convex or not too concave, and otherwise reduces it.

[35] A special case of this is that the shocks are independent and identically distributed.

[36] This does not follow from Proposition 4. In this case, the inverse relative marginal cost curve may be neither log convex nor log concave across the entire range of feasible quantities sold (and, given the symmetry, it cannot be log convex across the entire range).

Table 1: Impact of the efficient auction relative to the SSA in the symmetric case

	Demand curves (bidder preferences)		
	Log concave and concave	Log concave and convex	Log convex
Expected bidder surplus	Increases	Increases	Decreases
Expected auctioneer surplus	Depends on degree of concavity*	Increases	Increases

*The efficient auction increases expected auctioneer surplus relative to the SSA if the demand curves are not too concave, and otherwise reduces it.

Proposition 5, summarised in Table 1, partitions the space of demand curves into the subcases in which either only one side of the market benefits in expectation or the bidders and auctioneer share the welfare gain of the efficient auction.

Proposition 5 *In the symmetric case,*

(i) the efficient auction increases (decreases) expected bidder surplus relative to the SSA if both demand curves are log concave (log convex).

(ii) the efficient auction increases expected auctioneer surplus relative to the SSA if the demand curves are convex.

In the case in which the auctioneer's marginal cost of production is independent of the quantities produced of the other goods, $\frac{\partial^2 C(q_1,\ldots,q_J)}{\partial q_j \partial q_k} = 0 \ \forall j,k \in J, j \neq k$, this result extends to $J > 2$ goods.

Corollary 3 *In the symmetric case in which* $\frac{\partial^2 C(q_1,\ldots,q_J)}{\partial q_j \partial q_k} = 0 \ \forall j,k \in J, j \neq k$,

(i) the efficient auction increases (decreases) expected bidder surplus relative to the SSA if the demand curves are log concave (log convex).

(ii) the efficient auction increases expected auctioneer surplus relative to the SSA if the demand curves are convex.

5.2 Standard functional forms

I illustrate the results for the class of demand curves that are derived from the Generalised Pareto Distribution (GPD), which we can write as

$$d_j(p_j, \theta_j) = \left(1 + \frac{\xi_j(p_j - \theta_j - \mu_j)}{\sigma_j}\right)^{-\frac{1}{\xi_j}}, \ j = \{1, 2\}$$

where $\xi_j < 1$, $\sigma_j > 0$. Special cases of the GPD are below, specifically those that correspond to linear, log linear and constant elasticity demands.

I consider the cases in which the demand curves are both concave (and therefore log concave); both log concave but both convex; and both log convex (and therefore convex).[37]

1. If $\{\xi_1, \xi_2\} < -1$, the demand curves are strictly concave. If the inverse relative marginal cost curve, $\widetilde{mc}_1(\widetilde{p}_1)$, is also log convex (or $\widetilde{MC}_1(q_1) = 0$), the efficient auction reduces expected auctioneer surplus relative to the SSA (Propositions 3 and 4). Under these conditions, the efficient auction must increase expected bidder surplus.

   If $\frac{\mu_2}{\sigma_2} \geq \frac{1}{\xi_1}$, the expected price elasticity of demand for good 1 is smaller than that for good 2 at all allocations. Switching to the efficient auction also increases expected bidder surplus if $\{\xi_1, \xi_2\} < -1, \widetilde{MC}_1(q_1^S) \geq 0$ and $\frac{\mu_2}{\sigma_2} \geq \frac{1}{\xi_1}$ (Propositions 1 and 2).

2. If $-1 \leq \{\xi_1, \xi_2\} \leq 0$, the demand curves are log concave and convex. If, in addition, $\frac{\mu_2}{\sigma_2} \geq \frac{1}{\xi_1}$ and $\widetilde{MC}_1(q_1^S) \geq 0$, then the efficient auction increases expected bidder surplus relative to the SSA (Propositions 1 and 2), as described above. It increases expected auctioneer surplus if, in addition to the demand curves being convex, $\widetilde{mc}_1(\widetilde{p}_1)$ is log concave (or $\widetilde{MC}_1(q_1) = 0$) (Propositions 3 and 4).

3. If $0 \leq \{\xi_1, \xi_2\}$, the demand curves are log convex. If $\frac{1}{\xi_2} \leq \frac{\mu_2}{\sigma_2} \leq \frac{1}{\xi_1} \leq \frac{\mu_1}{\sigma_1}$, or, $\frac{\mu_2}{\sigma_2} \leq \frac{1}{\xi_2} \leq \frac{\mu_1}{\sigma_1} \leq \frac{1}{\xi_1}$, then the expected price elasticity of demand for good 1 is larger than that for good 2 at all allocations. If, in addition, $\widetilde{MC}_1(q_1^S) \leq 0$, then the efficient auction reduces expected bidder surplus relative to the SSA and must therefore increase expected auctioneer surplus.

   If, alternatively, $0 \leq \{\xi_1, \xi_2\}$ and $\widetilde{mc}_1(\widetilde{p}_1)$ is log concave, switching to the efficient auction increases expected auctioneer surplus.

I now show how the results of three standard cases of asymmetric goods correspond to those of the symmetric case.

Linear If $\xi_1 = \xi_2 = -1$, the demand for goods 1 and 2 are linear, so the Bidder Uncertainty Effect is positive and the Auctioneer Price Level Effect (with reference to either good) is zero. If the auctioneer's relative marginal cost curve with reference to good 2 is also linear, the Auctioneer Relative Price Effect with reference to good 2 is positive and, because the

[37] See Appendix F for details.

expected quantity sold of each good is the same in the two sales mechanisms, the Bidder Allocation Effect is zero. So in this case the efficient auction increases both expected bidder and auctioneer surplus relative to the SSA, as we anticipate from the symmetric case because the demand curves are log concave (so the bidders gain) but not strictly concave (so the auctioneer also unambiguously gains).[38]

Log linear As $(\xi_1, \xi_2) \to 0$, the demand curves become log linear, that is, $D_j(q_j, \theta_j) = \mu_j - \sigma_j log\,(q_j) + \theta_j, j = \{1, 2\}$. With log linear demand, uncertainty in the quantities sold has no impact on expected bidder surplus, so the Bidder Uncertainty Effect is zero. The impact of an increase in good j on the payment of inframarginal winners for good j is constant in its quantity sold, equal to σ_j. If $\sigma_1 \geq (\leq)\ \sigma_2$, the Bidder Allocation Effect is positive (negative), and the efficient auction increases (reduces) expected bidder surplus relative to the SSA.[39] The knife-edge with asymmetric goods between the cases in which $\sigma_1 > \sigma_2$ (the bidders gain) and $\sigma_1 < \sigma_2$ (the bidders lose) corresponds exactly to the knife-edge result of the bidders' indifference to the efficient auction in the symmetric case.

On the auctioneer side, the Auctioneer Price Level Effect (with reference to either good) is positive due to the log linear demand. If the auctioneer's inverse relative marginal cost curve, $\widetilde{mc}_1(\widetilde{p}_1)$, is log concave, including log linear, the Auctioneer Relative Price Effect with reference to good 2 is positive, so that the efficient auction increases expected auctioneer surplus relative to the SSA. And if the inverse relative marginal cost curve is not too log convex, the positive Auctioneer Price Level Effect outweighs the negative Auctioneer Relative Price Effect so that the auctioneer still benefits. The results for the auctioneer therefore correspond to the symmetric case if the auctioneer's marginal cost curve is either log concave or not too log convex.

Constant elasticity The demand curves have constant expected elasticities if $\xi_j = \frac{\sigma_j}{\mu_j}, \sigma_j > 0, \mu_j > 0, j = \{1, 2\}$. These demands are log convex so the Bidder Uncertainty Effect is neg-

[38]The linear model allows for direct analogy to the result in consumer theory that the surplus of a consumer with linear demand is determined by the mean and variance of the market price. In our setting, bidders' values for a good, and its auction price at a given allocation, change one-for-one with the common taste shock for that good. So expected bidder surplus for good j under mechanism m can be expressed in terms of the mean and variance of the auction price less the shock to good j, $(p_j^m - \theta_j)$. Conditional on the expected 'adjusted' price, expected bidder surplus is increasing in its variance.

[39]If $\sigma_1 > \sigma_2$, the demand for good 1 is relatively inelastic, so good 1's auction price falls by more than good 2's rises when an incremental unit is reallocated from good 2 to good 1, causing a greater change in bidder surplus for good 1. The efficient auction increases the expected quantity sold of the relatively profitable good.

ative. The Bidder Allocation Effect is also negative if $\widetilde{MC}_1(q_1^S) \leq 0$ and demand for good 1 is more elastic than for good 2, that is, if $\frac{\mu_1}{\sigma_1} \geq \frac{\mu_2}{\sigma_2}$. In this case, the efficient auction reduces expected bidder surplus and consequently must increase expected auctioneer surplus relative to the SSA. So under these conditions we have the same result as in the symmetric case that only the auctioneer benefits from the efficient auction.

The demand curves are convex, so the Auctioneer Price Level Effect with reference to either good is positive. So, exactly as in the log linear case, the efficient auction increases expected auctioneer surplus if the auctioneer's inverse relative marginal cost curve is not too log convex, again corresponding to the symmetric case.

5.3 Constant relative marginal costs

An important case is when the auctioneer's marginal cost of producing one good relative to another is constant. This arises if, for example, the marginal cost of each good is constant, or if the goods are interchangeable to produce, i.e., $C(x, 1-x) = C(1-x, x)\ \forall x \in [0, 1]$.

It follows automatically from the main analysis that the results of Propositions 1, 2 and 3 continue to hold exactly as before. Moreover it follows directly from its definition that the Auctioneer Relative Price Effect is zero. The reason is that it captures the impact on the auctioneer of variation in the relative price, but in the efficient auction the relative price equals the auctioneer's relative marginal cost so is always constant.

Proposition 6 *If the auctioneer's relative marginal cost is constant,*

(i) the efficient auction increases (decreases) expected bidder surplus relative to the SSA under the same conditions as in Corollary 1.

(ii) the efficient auction increases (decreases) expected auctioneer surplus relative to the SSA if the demand for good 2 (good 1) is convex (concave).

The empirical analysis in Section 7 is an example of this case.

6 Extensions

6.1 Profit-maximising auctions

The main analysis is most relevant to auctions held by public sector institutions, including the application studied in Section 7, in which efficiency is a key objective. In other contexts, it is more reasonable to assume that the auctioneer aims to maximise auctioneer surplus.

A PMA which maximises ex post auctioneer surplus is the "profit-maximising auction". The relevant comparator is the SSA in which the fixed quantities sold of the two goods are chosen to maximise expected auctioneer surplus. As in the main model, I assume the auctioneer always sells the entire unit.

The profit-maximising auction naturally increases expected auctioneer surplus. By defining the Bidder Uncertainty Effect and Bidder Allocation Effect in an analogous way to that in Section 3, it is straightforward to see that the Bidder Uncertainty Effect is positive (negative) if the demand curves are log concave (log convex), as in the main analysis. The objective is irrelevant to whether the bidders gain or lose as a group from the uncertainty in the quantities sold in the PMA.

I relabel the goods so that the profit-maximising auction increases the expected quantity sold of good 1.[40] The Bidder Allocation Effect is therefore positive (negative) if good 1 is relatively more (less) profitable on the margin than good 2 for the bidders. This is the case if reallocating a marginal unit to good 1 from good 2 increases the bidders' gross value by more (less) than it increases the total payment they make to the auctioneer, i.e., auctioneer revenue. But the relevant SSA equates the auctioneer's expected marginal surplus across the two goods: the expected marginal revenue of good 1 relative to good 2 equals the auctioneer's relative marginal cost. So the Bidder Allocation Effect is positive (negative) if the difference in the expected values of the marginal winners is greater (less) than the auctioneer's relative marginal cost.

Proposition 7 *Relative to the SSA in which the fixed shares sold of the two goods are chosen to maximise expected auctioneer surplus,*

(i) the profit-maximising auction increases expected auctioneer surplus.

(ii) the profit-maximising auction increases (decreases) expected bidder surplus if both demand curves are log concave (log convex),[41] *and the expected inverse demand for good 1 relative to good 2 is greater (less) than the relative marginal cost of good 1 at the quantities sold in the SSA.*[42]

[40] A sufficient condition for this labelling is that marginal revenue for good 1 relative to good 2 is weakly more convex than the relative marginal cost of good 1 (see Appendix B).

[41] Naturally, it is sufficient that the condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j,\theta_j) : s_j(p_j,\theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j = \{1,2\}$.

[42] As in Proposition 1, it is sufficient that the condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j,\theta_j) : d_j(p_j,\theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j = \{1,2\}$. The effect is strictly positive (negative) if at least one demand curve is strictly log concave (log convex).

It of course follows that, in the symmetric case, the conditions under which the profit-maximising auction increases (decreases) expected bidder surplus are the same as those under which the PMA increases (decreases) expected bidder surplus in the main model (described in Table 1).[43]

6.2 Variable total supply

The total supply available at auction is often fixed in practice.[44] However, a total supply which varies in response to the strength of demand is another source of efficiency gains.

I compare the welfare effects of an efficient auction with variable supply to an SSA with fixed quantities sold of the two goods, assuming that the goods are independent. The efficient auction is equivalent to running two separate simultaneous auctions with variable supply. This allows us to isolate the impact of the efficiency gains specifically arising from variable supply.

In the SSA (with fixed supplies), the quantity sold of each good is fixed ex ante to maximise expected welfare, and in the variable-supply efficient auction, the quantities sold of the goods maximise welfare. And so, $p_j(q_j^S, \theta_j) = D_j(q_j^S, \theta_j), \mathbb{E}\left[D_j(q_j^S, \theta_j)\right] = \mathbb{E}\left[\frac{\partial C(q_1^S, q_2^S)}{\partial q_j}\right]$ and $p_j(q_j^E(\boldsymbol{\theta}), \theta_j) = D_j(q_j^E(\boldsymbol{\theta}), \theta_j) = \frac{\partial C(q_1^E(\boldsymbol{\theta}), q_2^E(\boldsymbol{\theta}))}{\partial q_j}$. I consider the symmetric case of independent goods, that is, $\frac{\partial^2 C(q_1, \ldots, q_J)}{\partial q_j \partial q_k} = 0 \ \forall j, k \in J, j \neq k$, and variable supply.

In this case, the bidder and auctioneer sides of the market are analogous, and the impact on their surpluses can each be decomposed into the effects of the efficient auction on the uncertainty and expectation of the quantities sold. The gain to bidders from the uncertainty in the efficient auction does not depend on the total quantity sold, and as in Proposition 1, bidders benefit (lose) from the uncertainty of the quantities sold in the efficient auction if the demand curves are log concave (log convex). Similarly, the auctioneer benefits (loses) from this uncertainty if her inverse marginal cost curves are log concave (log convex).

[43] This is also the case for any objective function that treats the two goods symmetrically because the objective is then irrelevant to whether bidders benefit from the uncertainty in the quantities sold in the PMA. Fabra and Montero (2023) consider an objective function which is a weighted sum of welfare and revenue for the linear case.

[44] In the multi-product case, the central bank of Mexico permits a fixed nominal value of securities to be traded in its debt exchange auctions and the total quantity of energy procured in renewable energy auctions is often fixed (Szabó et al., 2020). The quantity sold in separate, simultaneous auctions of primary government debt is also typically fixed, including in Mexico's (see Section 7).

Nonetheless, there are exceptions. The quantity of liquidity issued in the Bank of England's ILTR operations was fixed until 2014, and now varies in response to the bids received and the Bank's privately known preferences. In the Federal Reserve's quantitative easing auctions, the total quantity sold was within a prespecified range.

The bidders and auctioneer both gain (lose) from an increase in the expected quantity sold of each good. With variable supply, the efficient auction increases (reduces) the expected quantity sold of each good relative to the SSA if the inverse demand curve for each good is more (less) convex than the marginal cost curve.

Proposition 8 *In the symmetric case of independent goods and variable supply,*

(i) the efficient auction increases (decreases) expected bidder surplus relative to the SSA if the inverse demand curve for each good is more (less) convex than the marginal cost curve, that is, $\frac{\partial^2 D_j(q_j,\theta_j)}{\partial q_j^2} \geq (\leq) \frac{\partial^2 MC_j(q_j)}{\partial q_j^2}, j = \{1,2\}$*, and the demand curves are log concave (log convex).*

(ii) the efficient auction increases (decreases) expected auctioneer surplus relative to the SSA if the inverse demand curve for each good is more (less) convex than the marginal cost curve and the inverse marginal cost curves are log concave (log convex).[45,46]

6.3 Alternative mechanism: reference price auction

The reference price auction (RPA) is a standard mechanism which only partially accommodates the imperfect substitutability of the two goods. Good 1 is assigned a fixed reference price (and the reference price of good 2 is normalised to zero). Within a multi-product auction, the total quantity to be sold is fixed at one unit and the acceptance priority of bids is determined by their differences with their respective reference prices. The auction clearing price of each good equals the marginal losing bid for that good.

The RPA is used in both financial and energy markets. The Bank of England and Federal Reserve used RPAs for their quantitative easing purchase schemes (Bank of England, 2022; Song and Zhu, 2018).[47] Both Germany and the UK use RPAs to allocate renewable energy contracts (Fabra and Montero, 2023).

The RPA fixes the relative price of the two goods. If the auctioneer's marginal cost of producing good 1 relative to good 2, $\widetilde{MC}_1(q_1)$, is constant, the optimal relative price is

[45] In the special case in which the auctioneer's relative marginal cost is constant (which violates the model assumption of strictly increasing marginal costs), the auctioneer is indifferent between the two designs.

[46] The effects are each strictly positive (negative) if at least one of their corresponding conditions holds strictly.

[47] Armantier, Holt and Plott (2013) report that the US Treasury also considered using an RPA to purchase "toxic assets" (namely mortgage-backed securities) in the $700 billion Troubled Asset Relief Program (TARP) in 2008.

Table 2: Impact of the efficient auction on expected surplus relative to the RPA in the symmetric case

	Demand curves (bidder preferences)		
	Log concave and concave	Log concave and convex	Log convex
Expected bidder surplus	Decreases	Decreases	Increases
Expected auctioneer surplus	Increases	Increases	Ambiguous

constant so the RPA is identical to the efficient auction. With more general auctioneer preferences, the RPA is inefficient.

In the symmetric case, the optimal reference price for good 1 is zero, and the RPA is equivalent to running one auction in which bids for the two goods are treated interchangeably. For each good, the expected quantity sold is equal across the two mechanisms and only the impact of uncertainty is relevant to their distributional effects.

The RPA creates *more* uncertainty in the quantities sold than the efficient auction. This is because the quantities sold of the two goods in the RPA adjust to the relative strength of demand, characterised by the bids, but not to the auctioneer's increasing relative marginal cost. So the conditions under which the bidders benefit from the efficient auction relative to the RPA are opposite to the conditions under which they benefit relative to the SSA. The efficient auction increases (decreases) expected bidder surplus if the demand curves are log convex (log concave), as the bidders gain (lose) from the reduced uncertainty.

Because the efficient auction increases the sum of bidder and auctioneer surplus, the efficient auction must increase expected auctioneer surplus if the demand curves are log concave.[48]

Proposition 9, summarised in Table 2, partitions the space of demand curves into the subcases in which one or both sides of the market benefit in expectation from the efficient auction relative to the RPA.

Proposition 9 *In the symmetric case,*

(i) the efficient auction increases (decreases) expected bidder surplus relative to the RPA

[48]For auctioneer surplus, we can separate the impact of the uncertainty in the relative prices of the two goods from the overall price level as in Section 4. However, only the second effect is straightforward: if the demand curves are concave (convex), the greater uncertainty in the quantities sold in the RPA reduces (increases) the expected price level relative to the efficient auction. So the auctioneer gains (loses) from the impact of the efficient auction on the expected price level.

if the demand curves are log convex (log concave).

(ii) the efficient auction increases expected auctioneer surplus relative to the RPA if the demand curves are log concave and has an ambiguous impact if the demand curves are log convex.[49]

7 Primary market for Mexican government debt

Primary auctions are currently held separately and simultaneously for different Mexican Treasury bonds. I estimate the effects of instead using efficient multi-product auctions to jointly allocate bonds of different maturities.[50]

To recover bidders' and the government's preferences over the bonds, I assume bidders' bids correspond to their valuations and use a revealed preference approach to identify the government's preferences from their issuance choices. Equipped with these preferences, I can estimate the differences in average bidder and auctioneer surpluses between the existing mechanism to allocate bonds and a counterfactual Product-Mix Auction.

7.1 Institutional details and data

I analyse the primary market for Mexican Federal Treasury Certificates (CETES). These are zero-coupon bonds with terms to maturity of typically 28, 91, 182 and 364 days, which were the main source of federal government financing in 1978 – 2017 and represented 25% of Mexican government debt in 2001 – 2017 (Cole et al., 2022). They are issued via separate sealed-bid uniform-price auctions, which are held simultaneously on a weekly schedule. The supply issuance of each bond is fixed so the auctions on any given day can be characterised as SSAs.

I collect bid-level data (including both winning and losing bids and bidder identifiers) for every auction from October 2017—the date on which the current uniform-price protocol was introduced—to August 2021. This constitutes 205 observations of the 28-, 91- and 182-day bond auctions, and 98 of the 364-day bond auctions.[51] To my knowledge, this is the first

[49] In the special case in which the auctioneer's relative marginal cost is constant (which violates the model assumption of strictly increasing marginal costs), expected bidder surplus and expected auctioneer surplus are each identical in the efficient auction and SSA.

[50] This is a particularly natural alternative as Banco de México currently runs other multi-product debt-exchange auctions. The quantity traded of each bond in "subastas de vasos comunicantes", as well as in "permutas de valores gubernamentales" (conducted on behalf of the Mexican federal government), depends on the bids submitted and the auctioneer's supply preferences.

[51] Auctions of all maturities were held weekly, except auctions of 364-day bonds are held only monthly in the period October 2017 – June 2020.

Table 3: Summary statistics of Mexican CETES primary auctions, 3 October 2017 – 31 August 2021

Maturity (days)		28	91	182	364
Number of auctions		205	205	205	98
Total quantity allocated	Mean	6,305	10,205	12,283	10,587
(million pesos)	Std. dev.	1,640	2,498	1,492	2,978
Number of bidders	Mean	20.37	20.28	16.24	15.16
	Std. dev.	3.42	2.89	3.09	2.74
Number of bids	Mean	64.80	53.14	45.22	42.80
	Std. dev.	12.23	10.84	11.01	12.08
Discount rate (weighted)	Mean	6.52	6.55	6.43	5.65
	Std. dev.	1.52	1.51	1.49	1.45
Bid price (weighted)	Mean	9.95	9.84	9.68	9.43
	Std. dev.	0.01	0.04	0.07	0.14
Cover ratio	Mean	3.25	3.07	2.68	2.67
	Std. dev.	0.67	0.67	0.64	0.61
Quantity demanded by bidder	Mean	15.95	15.14	16.49	17.63
(% of total allocated)	Std. dev.	18.06	17.77	16.88	17.08
Quantity allocated to winner	Mean	6.94	7.23	9.28	9.64
(% of total allocated)	Std. dev.	9.69	12.35	12.32	12.84

The cover ratio is the aggregate quantity demanded in an auction as a proportion of the total issuance. Discount rates and bid prices are weighted by the quantities demanded at those prices by bidders. The standard deviations of discount rate, bid price, quantity demanded by bidder and quantity allocated to winner are across auctions and bidders.

analysis of this period in which the uniform-price protocol was used.[52]

Bids Each bid for a particular bond specifies the quantity that the bidder demands and the discount rate that they are willing to accept. Bidders can submit multiple bids with minimum increments of one basis point and minimum quantity increments of 5,000 Mexican pesos (MEX$). I convert the yields into bid prices so that bid functions (describing the bidder's bid as a function of the quantity demanded) are decreasing in quantity.

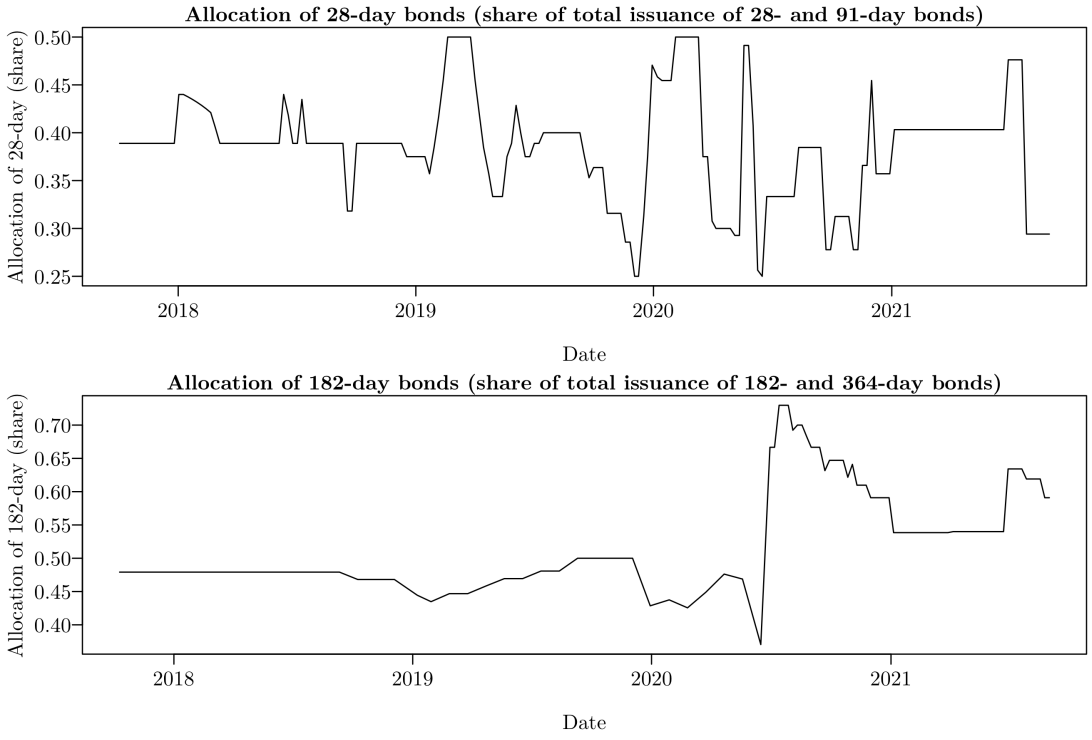
Descriptive statistics are provided in Table 3.

Issuance For each bond, the issuance is fixed at a quantity which is announced typically four days before the auction.[53] Both the total issuance of CETES and the share issued of each bond vary over time (Figure 4).

[52]A discriminatory-pricing rule was used from November 1995 to October 2017 and studied by Castellanos and Oviedo (2008) and Cole et al. (2021). Umlauf (1993) studied a previous (1990) change from a discriminatory- to a uniform-pricing rule.

[53]The auctions are typically held on Tuesdays and the amount to be issued of each bond is announced on the last day of the previous working week.

Figure 4: Supply issuance of in Mexican CETES primary auctions



The federal government's annual financing plan lays out its long-term debt portfolio optimisation strategy (Gobierno de México, 2022). In the plan, CETES are grouped into two pairs of bonds—28- and 91-day, and 182- and 364-day bonds—and the total issuance of each pair of CETES is determined within an optimisation model, but the allocation within pair is not. Motivated by this, my counterfactual analysis focuses on the optimal allocation within pair, abstracting from the choice of the total issuance for the pair.

The government's preferences between CETES are unobserved. In line with both the large literature on sovereign debt portfolio optimisation (e.g., Missale and Blanchard, 1994; Greenwood et al., 2015a, b) and the optimal portfolio models implemented in practice in Mexico (Gobierno de México, 2022), as well as in, e.g., Canada (Bolder and Deeley, 2011), Sweden (Bergstrom et al., 2002), Turkey (Balibek and Memis, 2012), and the UK (Pick and Anthony, 2006), I characterise the government's objective by a tradeoff between minimising the costs of issuance and maintaining a prudent risk profile.[54] Loosely speaking, shorter-term bonds are favourable from a financial perspective—bidders are willing to pay a higher price—but are unfavourable from a risk perspective because shorter-term bonds need to be rolled over

[54]I abstract from the possibility that the choice of issuance between CETES is determined by cash management objectives—the IMF (2021) reports that active cash management by the Mexican government through CETES issuance was limited during the sample period.

more frequently and this occurs at uncertain interest rates.

Specifically, I assume the government's total cost of issuing a bond is equal to its nominal debt service charges (reflecting its financial cost) less its Macaulay duration (reflecting its benefits in reducing rollover risk), weighted by an unobserved time-varying preference for duration.[55] For zero-coupon CETES, the nominal debt service charge is simply the bond's face value and its Macaulay duration is its term to maturity. The government's surplus is the difference between total revenue and total cost, and so on date t is

$$AS_t(q_{1t}, q_{2t}) = \sum_{j=\{1,2\}} \left[p_{jt}(q_{jt}) - v_j + \alpha_t \frac{t_j}{364} \right] q_{jt} \tag{1}$$

where, for bond $j \in \{1, 2\}$ on date t, q_{jt} is its quantity issued, $p_{jt}(q_{jt})$ is its auction price, v_j is its face value, and t_j is its term in days, and α_t is an unknown preference parameter.

This is a simplified version of the preference specification used to allocate bonds in Banco de México's debt-exchange auctions ("subastas de vasos comunicantes") in which multiple securities, including CETES, are exchanged within a multi-product auction. My interpretation of conversations with practitioners at Banco de México is that this is a reasonable approximation for the government's preferences between CETES.

7.2 Method

I group the bonds into two pairs: 28- and 91-day CETES, and 182- and 364-day CETES. For each pair, I compare the bidder and government surpluses under the actual mechanism to a counterfactual Product-Mix Auction (PMA) which jointly allocates the two bonds to maximise welfare (conventionally defined by the sum of total bidder and government surplus) given the same fixed total issuance.[56] The share issued of each of the two bonds is fixed in advance in the original mechanism but varies in the PMA. Surpluses are averaged over the set of dates T on which both bonds are auctioned. The following method refers to one pair of bonds, labelled bonds 1 and 2.

[55] Allen, Kastl and Wittwer (2022) use a similar specification to approximate the difference in the Bank of Canada's surplus from issuing Treasury bills of different maturities. They approximate the issuance cost of one bond relative to another by its initial relative price. More generally, alternative metrics of the cost and risks of issuance may be relevant to the choice over issuance of Treasury bills (see, e.g., Gobierno de México (2022) and Bolder and Deeley (2011)).

[56] In principle, we could run a counterfactual auction which jointly allocates all four bonds to maximise welfare. I focus on jointly allocating the bonds within pairs, as they are more plausibly substitutable so that Equation 1 is a more reasonable approximation—indeed, the government's annual financing plan allocates the bonds across but not within the pairs (Gobierno de México, 2022).

Estimating bidder preferences I assume that the submitted bids correspond to bidders' marginal valuations for the bonds, that bidders do not view the bonds as substitutes.[57] Each bidder submits a bid function, specifying the price they are willing to pay as a function of the quantity they demand. The bid function of bidder i for bond j in auction t is denoted $b_{ijt}(q_{ijt})$. I assume bidders are risk neutral so their surpluses are the differences between their bids and the auction prices. Bidder i's surplus for good j is the area between their bid function and the auction price for the quantity that they win:

$$\int_0^1 \mathbb{I}\big(b_{ijt}(q_{ijt}) \geq p_j(q_{jt})\big)\Big(b_{ijt}(q_{ijt}) - p_{jt}(q_{jt})\Big)dq_{ijt}$$

where $\mathbb{I}(.)$ is the indicator function; q_{jt} is the issuance of bond $j \in \{1,2\}$ on date t and $p_j(q_{jt})$ is its auction price. Total bidder surplus on date t, $BS_t(q_{1t}, q_{2t})$, is the sum of these surpluses over goods and winning bidders, i.e.,

$$BS_t(q_{1t}, q_{2t}) = \sum_{j=\{1,2\}} \sum_{i \in N_{jt}} \int_0^1 \mathbb{I}\big(b_{ijt}(q_{ijt}) \geq p_j(q_{jt})\big)\Big(b_{ijt}(q_{ijt}) - p_j(q_{jt})\Big)dq_{ijt} \tag{2}$$

where N_{jt} is the set of bidders for bond $j \in \{1,2\}$ on date $t \in \{1, ..., T\}$.

Estimating government preferences The government's (i.e., auctioneer's) surplus on date t is described by Equation 1. The preference parameters, $\alpha_t \ \forall t \in T$, are unobserved and are calibrated to match the share of the total issuance which is allocated of bond 1, q_{1t}^S (and corresponding of bond 2, $q_{2t}^S = (1 - q_{1t}^S)$). The observed choices are shown in Figure 4.

The calibration method is based on revealed preference: I assume that the allocation (q_{1t}^S, q_{2t}^S), which is fixed in advance of each auction, is chosen by the government to maximise expected welfare, i.e., the expected sum of auctioneer and bidder surplus defined in Equations 1 and 2.[58]

[57] Bidding according to their valuations is optimal if bidders are price takers so that they cannot exert unilateral market power by submitting bids below their valuations to push down the auction price. For the 28-day bond, for example, there are on average 20.4 bidders per auction, and the mean and standard deviation of the allocations of winning bidders are 6.9% and 9.7% of the total issuance, respectively, suggesting that this might be a reasonable approximation. Cole et al. (2022) also assume bidders behave competitively in their analysis of auctions of CETES in the period in which a discriminatory pricing rule was used, which precedes my sample period.

Allen, Kastl and Wittwer (2022) find that bills with 3-, 6-, and 12-month terms are weak substitutes for the average bidder in Canadian Treasury auctions.

[58] My interpretation of conversations with staff at Banco de México is that the government's objective for the auctions is not only to ensure issuance is cost-effective, but also to ensure well-functioning markets for the CETES, and that efficiency is an important objective. While positive weight might be placed on both auctioneer surplus and efficiency, I simplify the main analysis by assuming efficiency, i.e., maximising welfare, is the only objective. This allows for clear comparisons with the theoretical model. I also repeat the

In order to recover the preference parameters from the data, we must estimate the government's beliefs over the distribution of bids that will be submitted in auction t, given the information, denoted $I_{t,D}$, available at the time of the decision.

These beliefs and a particular value of the parameter, α_t, imply an optimal share to allocate of bond 1, $q_{1t}^S(\alpha_t)$, which is the share that maximises expected welfare over $q_{1t} \in [0, 1]$. The preference parameter, α_t, is recovered by matching the observed share to the optimal one.

To estimate the expected distribution of bids, I split the bid function of each potential bidder into two components. Bidder i's bid function for bond j in auction t, $b_{ijt}(q_{ijt})$, is the sum of the secondary market price at the time of the auction, s_{jt}, and a "residualised bid function", $\epsilon_{ijt}(q_{ijt})$, i.e., $b_{ijt}(q_{ijt}) = s_{jt} + \epsilon_{ijt}(q_{ijt})$. The residualised bid function captures the bidder's idiosyncratic valuation for the bond as well as factors which impact the valuation of all bidders which are not captured by the secondary market price. The expected bid function of bidder i for bond j in auction t at the time of the issuance decision is therefore

$$\mathbb{E}\left[b_{ijt}(q_{ijt})|I_{t,D}\right] = \mathbb{E}\left[s_{jt}|I_{t,D}\right] + \mathbb{E}\left[\epsilon_{ijt}(q_{ijt})|I_{t,D}\right]$$

where $\epsilon_{ijt}(q_{ijt})$ is the difference between the bidder's bid and the secondary market price, as a function of the quantity he demands.

I estimate the two components of the bid functions separately. First, the expected secondary market price on the day of auction t is estimated by the secondary market price on the day of the decision, which is typically five days prior to the auction.[59,60]

Second, to estimate the auctioneer's beliefs over the distribution of the residualised bids, I apply the resampling method developed by Hortaçsu (2002), Kastl (2011) and Hortaçsu and Kastl (2012). This approach is now standard in the empirical auction literature to estimate bidders' beliefs about the distribution of bids submitted by other bidders; I adapt it to simulate the auctioneer's beliefs.

Each bid function observed in the data represents one realisation of a bidder's marginal

exercise assuming the fixed shares maximise expected auctioneer surplus and compare the outcomes to the multi-product auction that maximises auctioneer surplus.

[59]The decision is made the day before the pair of auctions is announced. At this point, the federal government's decision is communicated to Banco de México, which runs the auction.

[60]This is an approximation, which is equivalent to assuming that the secondary market price follows a random walk. We expect that natural alternative forecasting techniques would produce similar results because data is only available at a daily frequency and the forecast is just five days ahead. I use the secondary yield provided by Valuación Operativa y Referencias de Mercado (Valmer) converted to price.

valuation function so the empirical distribution of residualised bids can be used to estimate the population distribution of residualised bids that the auctioneer expects to face in auction t.

To estimate the distribution of residualised bids in auction t, I resample from the the set of bid functions submitted within auction $t-1$, which is the last auction observed before the decision for auction t. Restricting the sample in this way avoids the bias that unobserved heterogeneity across auctions would cause if we pooled the bid functions across auctions. Moreover, the distribution of bid functions in the most recent auction will be most informative of the distribution in the next auction if there is a trend in the distribution over time.[61] Nonetheless, resampling bids from only one auction reduces the precision of the results and we do require that the bidders' valuation functions are identically distributed for the estimator to be consistent.[62]

For each auction t and bond $j \in \{1, 2\}$, there are N_{jt} potential bidders, each submitting a bid function. I draw N_{jt} bid functions observed in the data for bond j in auction $t-1$ with replacement (with bids equal to zero for non-participating bidders). This simulates one possible realisation of residualised bid functions, equal to the drawn bid functions less the secondary market price on the day of auction $t-1$. Repeating this resampling a large number of times simulates the full distribution of residualised bids the government expects to face for bond j in auction t.

The two estimated components—the estimated secondary market price and estimated distribution of residualised bids—provide an estimate of the government's beliefs over the distribution of bids in auction t at the time of the decision over the share to allocate of each bond. For any given α_t, the estimated optimal share allocated of bond 1, q_{1t}^S, is the share that maximises average welfare given the estimated distribution of bids. The parameter α_t is calibrated by matching the optimal share to the observed one.

Surpluses For each date $t \in T$, the counterfactual PMA selects the shares allocated to the two bonds, (q_{1t}^E, q_{2t}^E), which maximise welfare given the bids submitted on date t and given the calibrated α_t. The differences in surpluses between this counterfactual and the actual

[61] The residualised bid functions of a bidder depend on his information set, which includes both independent private signals and correlated signals which are not captured by the secondary market price. The presence of correlated signals would imply within-auction correlation of residualised bid functions. If the residualised bid functions were instead independent across bidders and their distribution was stable across auctions, we could use the full set of bid functions observed in the sample period.

[62] Cassola, Hortaçsu and Kastl (2013) discuss the consistency of the estimator. In Appendix G, I use an alternative resampling procedure, which relaxes this assumption.

mechanism are calculated given the bids submitted on date t. Estimated government and bidder surpluses in auction t, given an allocation (q_{1t}, q_{2t}), are the empirical counterparts to Equations 1 and 2 and are denoted $\hat{AS}_t(q_{1t}, q_{2t})$ and $\hat{BS}_t(q_{1t}, q_{2t})$.

Decomposition The Bidder Uncertainty Effect is estimated by

$$\frac{1}{T}\sum_{t\in T}\left(\hat{BS}_t(q^E_{1t}, q^E_{2t}) - \hat{BS}_t(\bar{q}^E_{1t}, \bar{q}^E_{2t})\right)$$

and the Bidder Allocation Effect by

$$\frac{1}{T}\sum_{t\in T}\left(\hat{BS}_t(\bar{q}^E_{1t}, \bar{q}^E_{2t}) - \hat{BS}_t(q^S_{1t}, q^S_{2t})\right)$$

where $(\bar{q}^E_{1t}, \bar{q}^E_{2t}) = (\frac{1}{T}\sum_{t\in T} q^E_{1t}, \frac{1}{T}\sum_{t\in T} q^E_{2t})$ (see Section 3).

As in Section 5.3, the auctioneer's marginal cost of good 1 relative to good 2 is constant, so the Auctioneer Relative Price Effect is set to zero, and the estimated Auctioneer Price Level Effect equals the average difference in auctioneer surplus:

$$\frac{1}{T}\sum_{t\in T}\left(\hat{AS}_t(q^E_{1t}, q^E_{2t}) - \hat{AS}_t(q^S_{1t}, q^S_{2t})\right)$$

7.3 Results

Results are shown in Table 4 for the case in which the objective is to maximise total welfare.[63] The results for the case in which the objective is to maximise auctioneer surplus are shown in Table 5. For each pair of bonds, the tables' upper panels show the issuance shares and bidder surplus in the SSA and the lower panels show the differences in issuance and surpluses between the counterfactual PMA and the SSA.

When the objective is to maximise total welfare, naturally the PMA is more efficient. For the pair of 182-day and 364-day bonds, the estimated welfare gain is 0.97 basis points of the issuance, i.e., MEX$2.27 million per auction. For the 28- and 91-day pair, the estimated welfare gain is 0.21 basis points (MEX$0.34m) (Table 4). Moreover, even when the objective is to maximise auctioneer surplus, the PMA still increases total welfare (Table 5).

When maximising total welfare, the PMA also increases both average bidder surplus and

[63] I estimate standard errors treating the government preference parameters, α_t $\forall t \in T$, as known as in Backus, Conlon and Sinkinson (2021) and Roussille and Scuderi (2022).

average auctioneer surplus for both pairs of bonds. For the 182-day and 364-day pair, the estimated bidder and auctioneer gains are 0.61 basis points of the issuance (or MEX$1.44 million) and 0.36 basis points (MEX$0.84m), respectively. For the 28-day and 91-day pair, the estimated gains are 0.08 basis points (MEX$0.13m) and 0.13 basis points (MEX$0.22m), respectively.

These results can be understood in terms of the predictions of the model. I label the 91- and 364-day bonds by bond 1 and 28- and 182-day bonds by bond 2 so that the counterfactual PMA increases the average allocation of bond 1 in each pair, in line with the model.

The significant, positive Bidder Uncertainty Effect for the 182- and 364-day pair can be understood by the fact that the bid functions are close to linear, and therefore log concave, in the neighbourhoods of the quantities issued, suggesting bidders benefit as a group from the uncertainty in the quantities allocated in the PMA (Proposition 1). For the 28- and 91-day pair, the bid functions are also close to linear, and therefore log concave, and the Bidder Uncertainty Effect is also positive, albeit imprecisely estimated.

Similarly, the significant, positive Bidder Allocation Effect for the 28- and 91-day pair can be understood by the fact that demand is more price elastic for bonds with shorter terms to maturity.[64] The price elasticity of demand of bond 1 (the 91-day bond) is therefore smaller in absolute value than that of bond 2's (28-day), and Proposition 2 implies that the Bidder Allocation Effect is positive if the relative marginal cost of good 1 is positive or not too negative at the quantities allocated in the SSA.[65] For the 182- and 364-day pair, our propositions do not predict the Bidder Allocation Effect (which is negative and insignificant).[66]

The impact on auctioneer surplus is explained by the impact on the price level alone as the auctioneer's constant marginal costs imply indifference to the uncertainty in relative prices. The PMA increases auctioneer surplus for the 28- and 91-day pair because it increases the price of the 28-day bond. The effect is positive but imprecisely estimated for the 182- and 364-day pair.

[64] In the neighbourhood of the SSA allocation, the elasticity of the 28-day bond (estimated to be $\hat{\epsilon}^D_{1,28} = -12964$) is larger in absolute value than the 91-day bond ($\hat{\epsilon}^D_{2,91} = -7508$), and the elasticity of the 182-day bond ($\hat{\epsilon}^D_{2,182} = -2718$) is larger than the 364-day bond ($\hat{\epsilon}^D_{1,364} = -1614$). Allen, Kastl and Wittwer (2022) also find that demand for bonds is more price elastic for shorter terms in Canadian Treasury bill auctions.

[65] The auctioneer's relative marginal cost of the 91-day bond (bond 1) is negative because its longer term is associated with less rollover risk, but the overall Bidder Allocation Effect is positive because bidders benefit from the larger expected quantity sold of the 91-day bond in the PMA.

[66] Demand for bond 1 (the 364-day bond) is less price elastic than demand for bond 2 (the 182-day bond) and the relative marginal cost of good 1 is sufficiently negative that Proposition 2 does not apply.

When the objective is to maximise auctioneer surplus, the PMA naturally benefits the auctioneer but average bidder surplus is lower in the PMA than the original design.

Overall, the welfare gains of the counterfactual PMA are moderate. Under our assumptions that the fixed allocation across the two bonds is determined jointly and optimally, the fact that the decision is made shortly before the auctions take place limits the additional welfare gains of determining the allocation within the auction itself.[67]

[67]In the Bank of England's liquidity auctions, Chapter 3 finds much larger potential welfare gains from the PMA relative to an SSA for which the Bank of England commits to the allocation across goods further in advance. In this case, the PMA increases welfare (conventionally defined by the sum of the bidders' and Bank of England's surpluses) by approximately 50% relative to the SSA.

Table 4: Surplus estimates and model predictions in Mexican CETES primary auctions given an objective to maximise total welfare

Pair of bonds		28- & 91-day	182- & 364-day
Bond 1	(term, days)	91	364
Bond 2	(term, days)	28	182
Issuance of bond 1 in SSA	(mean, % of total)	39.679	54.631
		(0.556)	(1.153)
	(std. dev., % of total)	6.971	9.366
		(1.479)	(1.670)
Bidder surplus in SSA	(bps of issuance)	0.750	2.838
		(0.038)	(0.246)
	(MEX$ million per auction)	1.213	6.636
		(0.059)	(0.611)
Difference between counterfactual PMA and actual SSA			
Δ Issuance of bond 1	(mean, % of total)	8.228***	4.284
		(1.920)	(3.216)
	(std. dev., % of total)	23.577***	25.689***
		(1.282)	(1.869)
Δ Welfare	(bps of issuance)	0.211***	0.972***
		(0.035)	(0.181)
	(MEX$ million per auction)	0.342***	2.273***
		(0.056)	(0.430)
Δ Bidder surplus	(bps of issuance)	0.078**	0.614*
		(0.028)	(0.243)
	(MEX$ million per auction)	0.127**	1.437*
		(0.046)	(0.571)
Δ Auctioneer surplus	(bps of issuance)	0.133***	0.358
		(0.031)	(0.238)
	(MEX$ million per auction)	0.215***	0.836
		(0.049)	(0.557)
Bidder Uncertainty Effect	(bps of issuance)	0.024	0.688**
		(0.034)	(0.240)
	(MEX$ million per auction)	0.038	1.610**
		(0.055)	(0.567)
Bidder Allocation Effect	(bps of issuance)	0.055**	-0.074
		(0.018)	(0.093)
	(MEX$ million per auction)	0.088**	-0.173
		(0.030)	(0.218)

Note: bps denote basis points. Surpluses measured in bps are weighted by total quantities issued and surpluses measured in MEX$ are unweighted. Standard errors are in parentheses. The level of welfare and auctioneer surplus in the actual SSA, which depend on the auctioneer's unobserved total benefit from issuing the total supply, are unidentified. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 5: Surplus estimates and model predictions in Mexican CETES primary auctions given an objective to maximise auctioneer surplus

Pair of bonds		28- & 91-day	182- & 364-day
Bond 1	(term, days)	91	364
Bond 2	(term, days)	28	182
Issuance of bond 1 in SSA	(mean, % of total)	39.223	55.455
		(0.639)	(1.283)
	(std. dev., % of total)	7.545	10.136
		(0.598)	(0.933)
Bidder surplus in SSA	(bps of issuance)	0.750	2.754
		(0.036)	(0.268)
	(MEX$ million per auction)	1.213	6.427
		(0.056)	(0.659)
Difference between counterfactual PMA and actual SSA			
Δ Issuance of bond 1	(mean, % of total)	10.676***	2.316
		(1.954)	(2.972)
	(std. dev., % of total)	23.235***	23.409***
		(1.138)	(1.869)
Δ Welfare	(bps of issuance)	0.243***	1.992***
		(0.039)	(0.475)
	(MEX$ million per auction)	0.393***	4.649***
		(0.064)	(1.127)
Δ Bidder surplus	(bps of issuance)	-0.146***	-0.254
		(0.025)	(0.204)
	(MEX$ million per auction)	-0.237***	-0.593
		(0.039)	(0.475)
Δ Auctioneer surplus	(bps of issuance)	0.389***	2.246***
		(0.033)	(0.383)
	(MEX$ million per auction)	0.630***	5.242***
		(0.054)	(0.909)

Note: bps denote basis points. Surpluses measured in bps are weighted by total quantities issued and surpluses measured in MEX$ are unweighted. Standard errors are in parentheses. The level of welfare and auctioneer surplus in the actual SSA, which depend on the auctioneer's unobserved total benefit from issuing the total supply, are unidentified. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

8 Conclusion

Even when markets are competitive, the welfare gains of a more efficient auction do not necessarily benefit both sides of the market. I provide conditions under which the efficient auction reduces expected bidder surplus, and conditions under which it reduces expected auctioneer surplus, relative to running independent auctions of substitute goods.

However, the conditions required for either side to lose seem extreme. If the model is symmetric in the goods, the efficient auction benefits both sides of the market if the aggregate bid functions are convex and log concave. Bidders only lose if the bid functions are very convex, and the auctioneer only loses if they are very concave. Moreover, the results are even clearer when the auctioneer's objective is to maximise auctioneer surplus, rather than total welfare: the auctioneer of course benefits from the optimal mechanism, and the bidders benefit if the aggregate bid functions are log concave. I show how the results extend if the model is asymmetric in the goods.

The mechanism that is efficient in my model—the Product-Mix Auction (PMA)—probably yields greater benefits relative to running independent auctions than my many-bidder model suggests. First, if the number of bidders is small, the PMA is likely to increase competition and efficiency by reducing large bidders' incentives to strategically bid below their values. Second, with a small number of bidders, idiosyncratic shocks to bidders' values provide an additional source of efficiency gain from the PMA.[68] Finally my model assumes each bidder is only interested in one of the goods. But if the bidders as well as the auctioneer view the goods as imperfect substitutes, the PMA in general further increases efficiency by allowing

[68]The effects on the distribution of surplus between bidders and the auctioneer are non-trivial. Lings (2013) considers the special case of a finite number of identical bidders for interchangeable goods. He finds that the efficient auction most likely increases expected auctioneer surplus, and increases expected bidder surplus if the survival function of the bidders' private values is log convex. This can be understood in terms of Bulow and Klemperer's (2012) model of rationing, which corresponds to running an SSA. (This contrasts to the case of a large number of bidders facing correlated shocks, described in Section 5.1, in which rationing bears similarities to the efficient auction.) Across the two goods, the bidders with above-median values always win in the efficient auction. Randomness in bidders' idiosyncratic values creates variation in the ordering of bidders along the total demand curve, which is constructed from the sum of demands across the two goods. With fixed quantities sold of the two goods, running an SSA is therefore rationing along the medium part of the total demand curve. (In both the efficient auction and SSA, the bidders with high value win. However, in the efficient auction, the highest of the medium value bidders are served (all those with values above the price at which total supply equals total demand), but in the SSA, only those with medium values who demand the good with the lower realised auction price are served. And so, it is as if the SSA is rationing bidders, with random allocation along the medium value part of the total demand curve.) And so, if the survival functions, and hence expected demand curves, of bidders are log concave (log convex), bidders gain (lose) from the splitting up of markets.

bidders to make bids that express their relative values for the different goods.[69] So my model does not consider all the potential benefits of the optimal mechanism; I leave analysis of the distribution of these additional benefits to future work.

The empirical work provides additional support for the view that combining auctions typically benefits both sides of the market. Both the multi-product auction of 28-day and 91-day Mexican Treasury bonds, and that of the 182-day and 364-day bonds, increase both average bidder and auctioneer surplus. This can be explained largely by the fact that the aggregate bid functions are log concave and typically convex in the relevant range.

In conclusion, the fact that both sides of the market are likely to gain, both in theory and in practice, is a strong argument for combining competitive auctions of substitute goods.

[69] In the general PMA (see Klemperer (2008), bidders can submit "paired bids", which express the quantity demanded by the bidder across both goods and a separate price for each of the two goods. The bidder wins at most the quantity he demands, and if he wins he is allocated the good which gives him the highest net surplus (under the assumption that his bids correspond to his valuations).

Appendices

A Proofs

A.1 Proposition 1

Suppose both demand curves are log concave. (The case of log convex demand curves is analogous.) For good $j \in \{1,2\}$, let $\widehat{D}_j(q_j) = -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j$, which is independent of θ_j because of the additivity of shocks. Evaluated at quantity $q_j = d_j(p_j,\theta_j)$,

$$\widehat{D}_j(d_j(p_j,\theta_j)) = -\frac{d_j(p_j,\theta_j)}{\frac{\partial d_j(p_j,\theta_j)}{\partial p_j}} \implies \frac{\partial \widehat{D}_j(q_j)}{\partial q_j} = \frac{\partial}{\partial p_j}\left(-\frac{d_j(p_j,\theta_j)}{\frac{\partial d_j(p_j,\theta_j)}{\partial p_j}}\right)\frac{\partial D_j(q_j,\theta_j)}{\partial q_j}$$

$$= \frac{\partial^2(log(d_j(p_j,\theta_j)))}{\partial p_j^2}\left(\frac{d_j(p_j,\theta_j)}{\frac{\partial d_j(p_j,\theta_j)}{\partial p_j}}\right)^2 \frac{\partial D_j(q_j,\theta_j)}{\partial q_j}$$

$$\implies \frac{\partial \widehat{D}_j(q_j)}{\partial q_j} \geq 0 \Leftrightarrow \frac{\partial^2(log(d_j(p_j,\theta_j))}{\partial p_j^2} \leq 0$$

because $\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} < 0$, and similarly for the strict inequalities. Applying Jensen's inequality to the function $\int_0^{q_j^E(\boldsymbol{\theta})} \widehat{D}_j(q_j)\, dq_j$ of $q_j^E(\boldsymbol{\theta})$,

$$\text{Bidder Uncertainty Effect} = \mathbb{E}\left[\sum_{j=\{1,2\}}\left(\int_{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}^{q_j^E(\boldsymbol{\theta})} \widehat{D}_j(q_j)\, dq_j\right)\right] \geq 0$$

A.2 Proposition 2

Observe that

$$\text{Bidder Allocation Effect} = \mathbb{E}\left[\sum_{j=\{1,2\}}\left(\int_{q_j^S}^{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]} -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j\, dq_j\right)\right] = \int_{q_1^S}^{\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]} X(q_1)\, dq_1$$

in which $X(q_1) = \left(-\frac{\partial D_1(q_1,\theta_1)}{\partial q_1} q_1 + \frac{\partial D_2(1-q_1,\theta_2)}{\partial q_2}(1-q_1)\right)$, because the auctioneer sells precisely one unit so that $q_2(q_1) = (1-q_1)$. Evaluating the integrand at the quantities sold in the

SSA and writing $\epsilon_j^D(q_j, \theta_j) = \frac{D_j(q_j,\theta_j)}{\frac{\partial D_j(q_j,\theta_j)}{\partial q_j} q_j}$ for the elasticity of demand for good j,

$$X(q_1^S) = -\frac{\mathbb{E}\left[D_1(q_1^S, \theta_1)\right]}{\mathbb{E}\left[\epsilon_1^D(q_1^S, \theta_1)\right]} + \frac{\mathbb{E}\left[D_2(1 - q_1^S, \theta_2)\right]}{\mathbb{E}\left[\epsilon_2^D(1 - q_1^S, \theta_2)\right]}$$

because the shocks are additive. The pair of quantities sold in the SSA is chosen to maximise expected welfare, so $\mathbb{E}\left[D_2(1 - q_1^S, \theta_2)\right] = \mathbb{E}\left[D_1(q_1^S, \theta_1)\right] - \widetilde{MC}_1(q_1^S)$, and therefore

$$X(q_1^S) = \mathbb{E}\left[D_1(q_1^S, \theta_1)\right] \left(\frac{1}{\mathbb{E}\left[\epsilon_2^D(1 - q_1^S, \theta_2)\right]} - \frac{1}{\mathbb{E}\left[\epsilon_1^D(q_1^S, \theta_1)\right]} \right) - \frac{\widetilde{MC}_1(q_1^S)}{\mathbb{E}\left[\epsilon_2^D(1 - q_1^S, \theta_2)\right]}$$

So if $\mathbb{E}\left[\epsilon_2^D(1 - q_1^S, \theta_2)\right] \leq \mathbb{E}\left[\epsilon_1^D(q_1^S, \theta_1)\right]$ and $\widetilde{MC}_1(q_1^S) \geq 0$, then $X(q_1^S) \geq 0$.

If the demand curves for goods $j = \{1, 2\}$ are log concave $\forall\{(p_j, \theta_j) : d_j(p_j, \theta_j) \in [\underline{q}_j, \bar{q}_j]\}$, then $\frac{\partial X(q_1)}{\partial q_1} \geq 0$ as shown in Proposition 1, and $q_1^S \leq \mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]$. Combining these three properties, it follows that the Bidder Allocation Effect is positive.

If instead the demand curves for goods $j = \{1, 2\}$ are log convex $\forall\{(p_j, \theta_j) : d_j(p_j, \theta_j) \in [\underline{q}_j, \bar{q}_j]\}$, so $\frac{\partial X(q_1)}{\partial q_1} \leq 0$ and $\mathbb{E}\left[\epsilon_2^D(1 - q_1^S, \theta_2)\right] \geq \mathbb{E}\left[\epsilon_1^D(q_1^S, \theta_1)\right]$, and $\widetilde{MC}_1(q_1^S) \leq 0$, so $X(q_1^S) \leq 0$, then the Bidder Allocation Effect is negative.

A.3 Lemma 2

I consider the case of good 1; the case of good 2 is analogous. In equilibrium, $p_1(q_1^E(\boldsymbol{\theta}),\theta_1) - p_2(q_2^E(\boldsymbol{\theta}),\theta_2) = \widetilde{MC}_1(q_1^E(\boldsymbol{\theta}))$ and $p_1(q_1^S,\theta_1) - p_2(q_2^S,\theta_2) = \widetilde{MC}_1(q_1^S) + \theta_1 - \theta_2$, and so

$$
\begin{aligned}
&\mathbb{E}\left[TAS(q_1^E(\boldsymbol{\theta}),\boldsymbol{\theta}) - TAS(q_1^S,\boldsymbol{\theta})\right] \\
=&\mathbb{E}\Bigg[\left(p_1(q_1^E(\boldsymbol{\theta}),\theta_1)q_1^E(\boldsymbol{\theta}) + p_2(q_2^E(\boldsymbol{\theta}),\theta_2)q_2^E(\boldsymbol{\theta}) - \int_0^{q_1^E(\boldsymbol{\theta})} \widetilde{MC}_1(q_1)\ dq_1 - C(0,1)\right) \\
&- \left(p_1(q_1^S,\theta_1)q_1^S + p_2(q_2^S,\theta_2)q_2^S - \int_0^{q_1^S} \widetilde{MC}_1(q_1)\ dq_1 - C(0,1)\right)\Bigg] \\
=&\mathbb{E}\Bigg[\left(\widetilde{MC}_1(q_1^E(\boldsymbol{\theta}))q_1^E(\boldsymbol{\theta}) - \int_0^{q_1^E(\boldsymbol{\theta})} \widetilde{MC}_1(q_1)\ dq_1\right) - \left(\widetilde{MC}_1(q_1^S)q_1^S - \int_0^{q_1^S} \widetilde{MC}_1(q_1)\ dq_1\right) \\
&+ \big(p_2(q_2^E(\boldsymbol{\theta}),\theta_2) - p_2(q_2^S,\theta_2)\big) - (\theta_1-\theta_2)q_1^S\Bigg] \\
=&\mathbb{E}\left[\int_{q_1^S}^{q_1^E(\boldsymbol{\theta})} \left(\widetilde{MR}_1(q_1) - \widetilde{MC}_1(q_1)\right)\ dq_1\right] + \mathbb{E}\left[p_2(q_2^E(\boldsymbol{\theta}),\theta_2) - p_2(q_2^S,\theta_2)\right]
\end{aligned}
$$

A.4 Proposition 3

Let $\bar{D}_j(q_j) = (D_j(q_j,\theta_j) - \theta_j)$, which is a function of q_j only because inverse demand for good j is additive in the shock θ_j. Then,

$$\mathbb{E}\left[p_j(q_j^E(\boldsymbol{\theta}),\theta_j) - p_j(q_j^S,\theta_j)\right] = \left(\mathbb{E}\left[\bar{D}_j(q_j^E(\boldsymbol{\theta}))\right] - \bar{D}_j(\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right])\right) + \left(\bar{D}_j(\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]) - \bar{D}_j(q_j^S)\right)$$

If $d_j(p_j,\theta_j)$ is convex in p_j then, because demand for good j is strictly decreasing, $\bar{D}_j(q_j)$ is convex and $\bar{D}_j(\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]) \leq \mathbb{E}\left[\bar{D}_j(q_j^E(\boldsymbol{\theta}))\right]$. The converse is analogous.

By the labelling of the two goods, $q_1^S \leq \mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]$ and correspondingly $q_2^S \geq \mathbb{E}\left[q_2^E(\boldsymbol{\theta})\right]$, so $\bar{D}_2(q_2^S) \leq \bar{D}_2(\mathbb{E}\left[q_2^E(\boldsymbol{\theta})\right]) \leq \mathbb{E}\left[\bar{D}_2(q_2^E(\boldsymbol{\theta}))\right]$ if demand for good 2 is convex, and $\bar{D}_1(q_1^S) \geq \bar{D}_1(\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]) \geq \mathbb{E}\left[\bar{D}_1(q_1^E(\boldsymbol{\theta}))\right]$ if demand for good 1 is concave.

A.5 Proposition 4

Suppose the inverse relative marginal cost curve with reference to good $j = \{1, 2\}$ is log concave. Evaluated at quantity $q_j = \widetilde{mc}_j(\widetilde{p}_j)$,

$$\widetilde{MC}_j(\widetilde{mc}_j(\widetilde{p}_j)) = \frac{\widetilde{mc}_j(\widetilde{p}_j)}{\frac{\partial \widetilde{mc}_j(\widetilde{p}_j)}{\partial \widetilde{p}_j}} \implies \frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j} = \frac{\partial}{\partial \widetilde{p}_j}\left(\frac{\widetilde{mc}_j(\widetilde{p}_j)}{\frac{\partial \widetilde{mc}_j(\widetilde{p}_j)}{\partial \widetilde{p}_j}}\right)\frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j}$$

$$= -\frac{\partial^2(log(d_j(\widetilde{p}_j, \theta_j)))}{\partial \widetilde{p}_j^2}\left(\frac{\widetilde{mc}_j(p_j)}{\frac{\partial \widetilde{mc}_j(p_j)}{\partial p_j}}\right)^2 \frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j}$$

$$\implies \frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j} \geq 0 \Leftrightarrow \frac{\partial^2(log(\widetilde{mc}_j(\widetilde{p}_j))}{\partial \widetilde{p}_j^2} \leq 0$$

because $\frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j} > 0$, and similarly for the strict inequalities.

Let $h_j(q_j) = \widetilde{MR}_j(q_j) - \widetilde{MC}_j(q_j) = \frac{\partial \widetilde{MC}_j(q_j)}{\partial q_j} q_j$. Relative marginal cost is increasing, so that $h_j(q_j) \geq 0 \ \forall q_j$. By the labelling of the two goods, $\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right] \geq q_1^S$, so

$$\mathbb{E}\left[\int_{q_1^S}^{\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]} h_1(q_1)\ dq_1\right] \geq 0,\ \mathbb{E}\left[\int_{q_2^S}^{\mathbb{E}\left[q_2^E(\boldsymbol{\theta})\right]} h_2(q_2)\ dq_2\right] \leq 0$$

Applying Jensen's inequality to the function $\int_0^{q_1^E(\boldsymbol{\theta})} h_1(q_1)\ dq_1$, the Auctioneer Relative Price Effect with reference to good 2 is positive if $\widetilde{MC}_1(q_1)$ is log concave. Analogously, the Auctioneer Relative Price Effect with reference to good 1 is negative if $\widetilde{MC}_2(q_2)$ is log convex.

Of course, if the inverse relative marginal cost curve of good $j \in \{1, 2\}$ has zero slope, then $h_j(q_j) = 0 \ \forall q_j$ and the Auctioneer Relative Price Effect with reference to good j is zero.

A.6 Proposition 5

The two shocks, (θ_1, θ_2), are exchangeable with zero mean, and so the relative shock to good 1, $(\theta_1 - \theta_2)$, is symmetrically distributed around zero. The auctioneer's and bidders' preferences are each symmetric in the two goods. Using the equilibrium condition for the efficient auction, it follows that $q_j^E(\theta), j = \{1, 2\}$, is also symmetrically distributed, with $q_j^S = \mathbb{E}\left[q_j^E(\theta)\right] = \frac{1}{2}$.

Bidder surplus The conditions follow from Proposition 1 and the fact that the expected quantities sold in the two mechanisms are equal.

Auctioneer surplus Proposition 3 determine the conditions under which the Auctioneer Price Level Effect is positive/negative. To determine the sign of the Auctioneer Relative Price Effect, I suppress the dependence on the vector of shocks, and denote the distribution of $q_k^E, k \in \{1,2\}$, by $G_k(q_k^E)$. By the symmetry, $\widetilde{MC}_k(1-q_k^E) = -\widetilde{MC}_k(q_k^E)$, and so

$$\begin{aligned}
\mathbb{E}\left[\int_{\mathbb{E}\left[q_k^E\right]}^{q_k^E}\left(\widetilde{MR}_k(q_k)-\widetilde{MC}_k(q_k)\right)dq_k\right] &= \mathbb{E}\left[\widetilde{MC}_k(q_k^E)q_k^E - \int_{\frac{1}{2}}^{q_k^E}\widetilde{MC}_k(q_k)dq_k\right] \\
&= 2\int_{\frac{1}{2}}^{1}\left(\widetilde{MC}_k(q_k^E)\left(q_k^E-\frac{1}{2}\right)-\int_{\frac{1}{2}}^{q_k^E}\widetilde{MC}_k(q_k)dq_k\right)dG_k(q_k^E) \\
&= 2\int_{\frac{1}{2}}^{1}\left(\int_{\frac{1}{2}}^{q_k^E}\left(\widetilde{MC}_k(q_k^E)-\widetilde{MC}_k(q_k)\right)dq_k\right)dG_k(q_k^E) \\
&\geq 0
\end{aligned}$$

because $\frac{\partial\widetilde{MC}_k(q_k)}{\partial q_k} \geq 0 \ \forall q_k$.

A.7 Corollary 3

For each of the J goods, the same quantity will be sold in the SSA and PMA in expectation, $q_k^S = \mathbb{E}\left[q_k^E(\boldsymbol{\theta})\right] = \frac{1}{J}, k \in J$.

Bidder surplus Lemma 1 and Proposition 1 directly generalise to J goods, and the Bidder Allocation Effect is zero.

Auctioneer surplus Let $MC_k(q_k) = \frac{\partial C(q_1,\ldots,q_J)}{\partial q_k}, k \in J$, and $MC_k(q) = MC(q) \ \forall k, q$,

then, for $j \in J$,

$$\begin{aligned}
&\mathbb{E}\left[TAS(q_1^E(\boldsymbol{\theta}), q_2^E(\boldsymbol{\theta}), ..., q_{J-1}^E(\boldsymbol{\theta}), \boldsymbol{\theta}) - TAS(q_1^S, q_2^S, ..., q_{J-1}^S, \boldsymbol{\theta})\right] \\
=&\mathbb{E}\left[\sum_{k=1}^{J}\left(p_k(q_k^E(\boldsymbol{\theta}), \theta_k)q_k^E(\boldsymbol{\theta}) - p_k(q_k^S, \theta_k)q_k^S - \int_{q_k^S}^{q_k^E(\boldsymbol{\theta})} MC(q_k)\ dq_k\right)\right] \\
=&\sum_{k=1}^{J}\mathbb{E}\left[MC(q_k^E(\boldsymbol{\theta}))q_k^E(\boldsymbol{\theta}) - MC(q_j^E(\boldsymbol{\theta}))\left(\frac{1}{J}\right) - \int_{\frac{1}{J}}^{q_k^E(\boldsymbol{\theta})} MC(q_k)\ dq_k\right] \\
&+\mathbb{E}\left[p_j(q_j^E(\boldsymbol{\theta}), \theta_j) - p_j\left(\frac{1}{J}, \theta_j\right)\right] \\
=&J\mathbb{E}\left[\int_{\frac{1}{J}}^{q_j^E(\boldsymbol{\theta})}\left(MC(q_j^E(\boldsymbol{\theta})) - MC(q_j)\right)dq_j\right] + \mathbb{E}\left[p_j(q_j^E(\boldsymbol{\theta}), \theta_j) - p_j\left(\frac{1}{J}, \theta_j\right)\right] \qquad (3)
\end{aligned}$$

because $p_k(q_k^E(\boldsymbol{\theta}), \theta_k) - p_j(q_j^E(\boldsymbol{\theta}), \theta_j) = MC_k(q_k^E(\boldsymbol{\theta})) - MC_j(q_j^E(\boldsymbol{\theta}))$ and $p_k(q_k^S, \theta_k) - p_j(q_j^S, \theta_j) = MC_k(q_k^S) - MC_j(q_j^S) + \theta_k - \theta_j\ \forall j, k \in J$.

Because $\frac{\partial MC(q_j)}{\partial q_j} \geq 0$, $\int_{\frac{1}{J}}^{q_j^E(\boldsymbol{\theta})}\left(MC(q_j^E(\boldsymbol{\theta})) - MC(q_j)\right)dq_j \geq 0\ \forall \boldsymbol{\theta}$ and the first term in Equation 3 is positive. The sign of the second term follows from Proposition 3.

A.8 Proposition 7

The impact on expected auctioneer surplus is immediate. It is clear from Proposition 1 that the Bidder Uncertainty Effect in this case, analogous to that defined in Section 3, is positive (negative) if the demand curves are log concave (log convex).

Turning to the Bidder Allocation Effect in this case, let $X(q_1) = (D_1(q_1, \theta_1) - MR_1(q_1, \theta_1)) - (D_2(1 - q_1, \theta_2) - MR_2(1 - q_1, \theta_2))$. Then,

$$\begin{aligned}
X(q_1^S) =&\mathbb{E}\left[D_1(q_1^S, \theta_1)\right] - \mathbb{E}\left[D_2(1 - q_1^S, \theta_2)\right] - \mathbb{E}\left[MR_1(q_1^S, \theta_1)\right] - \mathbb{E}\left[MR_2(1 - q_1^S, \theta_2)\right] \\
=&\mathbb{E}\left[D_1(q_1^S, \theta_1)\right] - \mathbb{E}\left[D_2(1 - q_1^S, \theta_2)\right] - \widetilde{MC}_1(q_1^S)
\end{aligned}$$

where the last line follows from the quantities sold in the SSA maximising expected auctioneer surplus. And so, $X(q_1^S) \geq 0$ if $\mathbb{E}\left[D_1(q_1^S, \theta_1)\right] - \mathbb{E}\left[D_2(1 - q_1^S, \theta_2)\right] - \mathbb{E}\left[MR_1(q_1^S, \theta_1)\right] \geq \widetilde{MC}_1(q_1^S)$. The rest of the proof is identical to that of Proposition 2.

A.9 Proposition 8

Define the Bidder Uncertainty and Allocation Effects as in Section 3 and define the Auctioneer Uncertainty and Allocation Effects similarly,

$$\text{Auctioneer Uncertainty Effect} = \mathbb{E}\left[\sum_{j=\{1,2\}}\left(\int_{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}^{q_j^E(\boldsymbol{\theta})}\left(MAR_j(q_j)-MC_j(q_j)\right)dq_j\right)\right]$$

and

$$\text{Auctioneer Allocation Effect} = \mathbb{E}\left[\sum_{j=\{1,2\}}\left(\int_{q_j^S}^{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}\left(MAR_j(q_j)-MC_j(q_j)\right)dq_j\right)\right]$$

where $MC_j(q_j) = \frac{\partial C(q_1,q_2)}{\partial q_j}$ and $MAR_j(q_j) = \frac{\partial(MC_j(q_j)q_j)}{\partial q_j}$. Define the inverse marginal cost curve for good j, equal to $mc_j(p_j)$.

It follows immediately that the difference in expected total bidder (auctioneer) surplus between the efficient auction and SSA, $\mathbb{E}\left[TBS(q_1^E(\boldsymbol{\theta}),\boldsymbol{\theta}) - TBS(q_1^S,\boldsymbol{\theta})\right]$ ($\mathbb{E}\left[TAS(q_1^E(\boldsymbol{\theta}),\boldsymbol{\theta}) - TAS(q_1^S,\boldsymbol{\theta})\right]$), equals the sum of the Bidder (Auctioneer) Uncertainty Effect and the Bidder (Auctioneer) Allocation Effect.

It follows from analogy to Proposition 1 that the Auctioneer Uncertainty Effect is positive (negative) if both inverse marginal cost curves are log concave (log convex).[70]

Now consider the Bidder Allocation Effect and Auctioneer Allocation Effects. For $j = \{1,2\}$, let $h_j(q_j) = D_j(q_j,\theta_j) - \theta_j - MC_j(q_j)$, which is a function of q_j only, because inverse demand for good j is additive in the shock θ_j and the goods are independently produced. By the equilibrium conditions, $h_j(q_j^S) = 0$, and $h_j(q_j^E(\boldsymbol{\theta})) = -\theta_j$ so that $\mathbb{E}\left[h_j(q_j^E(\boldsymbol{\theta}))\right] = \mathbb{E}\left[-\theta_j\right] = 0$. If the inverse demand for good j is more convex than its marginal cost, then $h_j(q_j)$ is strictly decreasing and convex, so $h_j\left(\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]\right) \leq 0$ and $\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right] \geq q_j^S$. Conversely, if the inverse demand for good j is less convex than its marginal cost, then $\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right] \leq q_j^S$.

Both $(D_j(q_j,\theta_j) - MR_j(q_j,\theta_j)) = -\frac{\partial D_j(q_j,\theta_j)}{\partial q_j}q_j \geq 0$ and $(MAR_j(q_j,\theta_j) - MC_j(q_j)) = \frac{\partial MC_j(q_j)}{\partial q_j}q_j \geq 0$, and so the result follows.

[70] It is sufficient that the condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j,\theta_j) : d_j(p_j,\theta_j) \in [\underline{q}_j, \bar{q}_j]\}, j = \{1,2\}$. The effect is strictly positive (negative) if at least one demand curve is strictly log concave (log convex).

A.10 Proposition 9

Bidder surplus Let q_1^R be the quantity sold of good 1 in the RPA; $\mathbb{E}(q_1^E(\boldsymbol{\theta})) = \mathbb{E}(q_1^R(\boldsymbol{\theta}))$, so

$$\begin{aligned}
&\mathbb{E}\big[TBS(q_1^E(\boldsymbol{\theta}),\boldsymbol{\theta}) - TBS(q_1^R(\boldsymbol{\theta}),\boldsymbol{\theta})\big] \\
&\quad = \mathbb{E}\left[\sum_{j=\{1,2\}}\left(\int_{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}^{q_j^E(\boldsymbol{\theta})} (D_j(q_j,\theta_j) - MR_j(q_j,\theta_j))\,dq_j - \int_{\mathbb{E}\left[q_j^R(\boldsymbol{\theta})\right]}^{q_j^R(\boldsymbol{\theta})} (D_j(q_j,\theta_j) - MR_j(q_j,\theta_j))\,dq_j\right)\right]
\end{aligned}$$

Let $\widetilde{\boldsymbol{\theta}} = \theta_2 - \theta_1$ be the relative shock to good 1. If $\widetilde{\boldsymbol{\theta}} \in [0, \bar{\theta}_2 - \underline{\theta}_1]$, $q_1^E((\theta_1, \theta_1 + \widetilde{\boldsymbol{\theta}})) < \frac{1}{2}\ \forall \theta_1 \in [\underline{\theta}_1, \bar{\theta}_1]$, and $\widetilde{MC}_1'(q_1) \geq 0$, then the quantities sold in the PMA and RPA satisfy

$$D_1(q_1^E((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})),\theta_1) - D_2(1-q_1^E((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})),\theta_1+\widetilde{\boldsymbol{\theta}}) = \widetilde{MC}_1(q_1^E((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}}))) \leq \widetilde{MC}_1\left(\frac{1}{2}\right) = 0$$

$$\begin{aligned}
&\text{and } 0 = D_1(q_1^R((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})),\theta_1) - D_2(1-q_1^R((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})),\theta_1+\widetilde{\boldsymbol{\theta}}) \\
&\implies q_1^R((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})) \leq q_1^E((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})) \leq \frac{1}{2} \quad \forall\big\{\theta_1 \in [\underline{\theta}_1,\bar{\theta}_1], \widetilde{\boldsymbol{\theta}} \in [0,\bar{\theta}_2-\underline{\theta}_1]\big\}
\end{aligned}$$

and, analogously, $\frac{1}{2} \leq q_1^E((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})) \leq q_1^R((\theta_1,\theta_1+\widetilde{\boldsymbol{\theta}})) \quad \forall\big\{\theta_1 \in [\underline{\theta}_1,\bar{\theta}_1], \widetilde{\boldsymbol{\theta}} \in [\underline{\theta}_2 - \bar{\theta}_1, 0]\big\}$. So $q_1^E(\boldsymbol{\theta})$ second order stochastically dominates $q_1^R(\boldsymbol{\theta})$. By Proposition 1, if the demand curves are log convex (log concave), $(D_j(q_j,\theta_j) - MR_j(q_j,\theta_j)) \leq (\geq)\ 0$, so that $\mathbb{E}\left[TBS(q_1^E(\boldsymbol{\theta}),\boldsymbol{\theta}) - TBS(q_1^R(\boldsymbol{\theta}),\boldsymbol{\theta})\right] \geq (\leq)\ 0$.

Auctioneer surplus The efficient auction increases total expected surplus, so it increases expected auctioneer surplus if it reduces expected bidder surplus.

B Labelling the goods

Lemma 3 *If* $\frac{\partial^2(D_1(q_1,\theta_1)-D_2(1-q_1,\theta_2))}{\partial q_1^2} \geq \frac{\partial^2 \widetilde{MC}_1(q_1)}{\partial q_1^2}$ $\forall q_1 \in [\underline{q}_1, \bar{q}_1]$*, the efficient auction increases the expected quantity sold of good 1 relative to the SSA.*[71]

Proof. Let $h(q_1) = D_1(q_1,\theta_1) - \theta_1 - D_2(1-q_1,\theta_2) + \theta_2 - \widetilde{MC}_1(q_1)$, which is a function of q_1 only, because inverse demand for good $j = \{1,2\}$ is additive in the shock θ_j. By the equilibrium conditions, $h(q_1^S) = 0$, and $h(q_1^E(\boldsymbol{\theta})) = \theta_2 - \theta_1$ so that $\mathbb{E}\left[h(q_1^E(\boldsymbol{\theta}))\right] = \mathbb{E}\left[\theta_2 - \theta_1\right] = 0$. The function, $h(q_1)$, is strictly decreasing and convex, so $h\left(\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right]\right) \leq 0$, and it is also decreasing, so $\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right] \geq q_1^S$. ■

If the PMA instead maximises auctioneer surplus, as in the profit-maximising auction, the quantity sold of good 1 (and implicitly, of good 2) adjusts so that the marginal revenue of good 1 relative to good 2 equals the auctioneer's relative marginal cost. This implies the following analogous sufficient condition for labelling the goods.

Lemma 4 *If* $\frac{\partial^2(MR_1(q_1,\theta_1)-MR_2(1-q_1,\theta_2))}{\partial q_1^2} \geq \frac{\partial^2 \widetilde{MC}_1(q_1)}{\partial q_1^2}$ $\forall q_1 \in [\underline{q}_1, \bar{q}_1]$ *and* $\frac{\partial(MR_j(q_j,\theta_j)}{\partial q_j} < 0$ *for* $j = \{1,2\}$*, the profit-maximising auction increases the expected quantity sold of good 1 relative to the SSA in which the fixed shares sold of the two goods are chosen to maximise expected auctioneer surplus.*[72]

Proof. Let q_1^P be the quantity sold of good 1 in the profit-maximising auction and and q_1^T be the quantity sold of good 1 in the SSA in which the fixed shares sold of the two goods are chosen to maximise expected auctioneer surplus, respectively. Let $g(q_1) = MR_1(q_1,\theta_1) - \theta_1 - MR_2(1-q_1,\theta_2) + \theta_2 - \widetilde{MC}_1(q_1)$. In equilibrium, $g(q_1^P(\boldsymbol{\theta})) = \theta_2 - \theta_1$ and $g(q_1^T) = 0$. The rest of the proof is the same as that of Lemma 3, and so $\mathbb{E}\left[q_1^P(\boldsymbol{\theta})\right] \geq q_1^T$. ■

C Distribution of bidder surplus across markets

The impact of the efficient auction on the expected surplus of the bidders for a single good can be decomposed into good-specific Bidder Uncertainty and Allocation Effects. For good j, let the Bidder Uncertainty Effect$_j$ $= \mathbb{E}\left[\int_{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]}^{q_j^E(\boldsymbol{\theta})} \left(D_j(q_j,\theta_j) - MR_j(q_j,\theta_j)\right)\, dq_j\right]$ and the

[71] Of course, the implication also holds strictly, that is, if $\frac{\partial^2(D_1(q_1,\theta_1)-D_2(1-q_1,\theta_2))}{\partial q_1^2} > \frac{\partial^2 \widetilde{MC}_1(q_1)}{\partial q_1^2}$ $\forall q_1 \in [\underline{q}_1, \bar{q}_1]$, then $\mathbb{E}\left[q_1^E(\boldsymbol{\theta})\right] > q_1^S$.

[72] The implication also holds strictly.

Bidder Allocation Effect$_j$ $= \int_{q_j^S}^{\mathbb{E}\left[q_j^E(\boldsymbol{\theta})\right]} \left(D_j(q_j,\theta_j) - MR_j(q_j,\theta_j)\right)\ dq_j$. The following decomposition therefore immediately follows.

Lemma 5 *The difference between the efficient auction and SSA in expected surplus of bidders for good $j \in \{1,2\}$, $\mathbb{E}\left[TBS_j(q_j^E(\boldsymbol{\theta}),\theta_j) - TBS_j(q_j^S,\theta_j)\right]$, equals the sum of the good-specific Bidder Uncertainty Effect and Bidder Allocation Effect for good j.*

The Bidder Uncertainty Effect is the sum of the good-specific Bidder Uncertainty Effects, capturing the gain to bidders from the uncertainty in the quantity sold of each good (see Proposition 1).

Corollary 4 *Bidder Uncertainty Effect$_j$, $j \in \{1,2\}$, is positive (negative) if the demand curve for good j is log concave (log convex).*[73]

By the labelling of the two goods, the efficient auction increases the expected quantity sold of good 1, and so Bidder Allocation Effect$_1 \geq 0$ and Bidder Allocation Effect$_2 \leq 0$.

D Decomposition of auctioneer surplus

D.1 Relationship between Figures 2 and 3

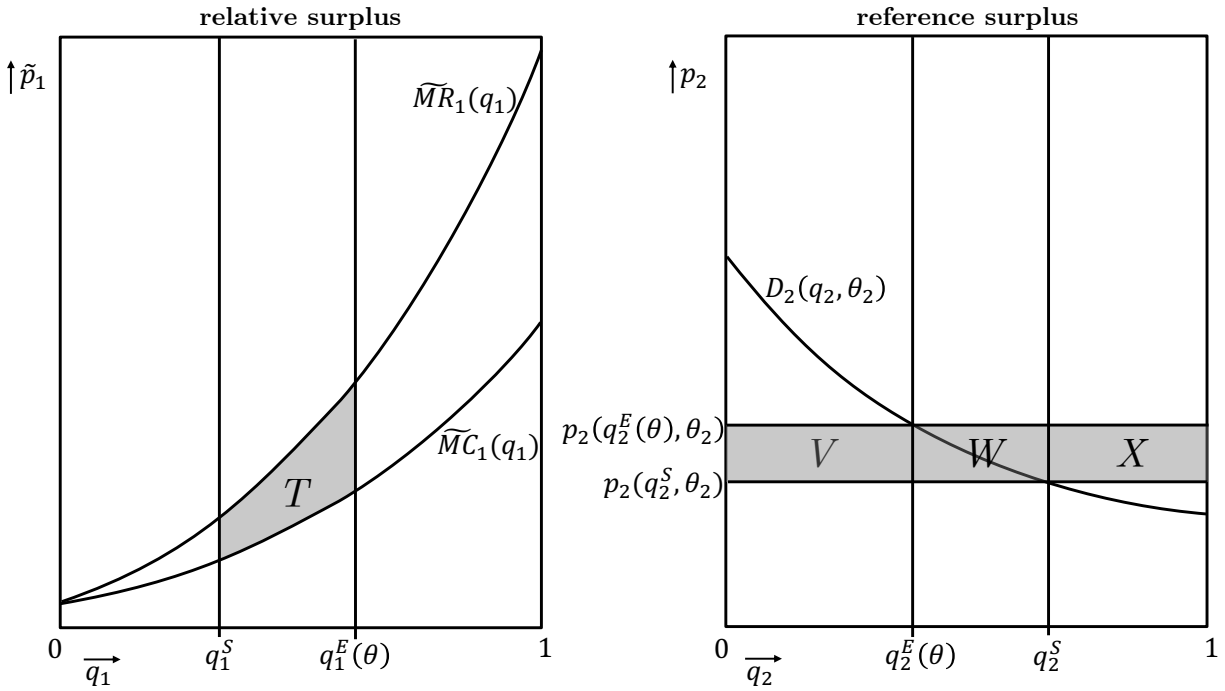
If the auctioneer has independent marginal costs of production, the areas shaded in Figure 2 correspond to those in Figure 3.

To see the relationship, I denote $\chi(\boldsymbol{\theta}) = (\tilde{MC}_1(q_1^E(\boldsymbol{\theta})) - \tilde{MC}_1(q_1^S))q_1^S$ and I split Area U in Figure 3 into three parts, Areas V, W and X, as illustrated in Figure 5. Comparing Figures 2 and 5, we have

$$\begin{aligned} \text{Area } T + \text{Area } W &= \text{Area } Q - \text{Area } S + \chi(\boldsymbol{\theta}) \\ \text{Area } V &= \text{Area } R \\ \text{Area } X &= (\theta_1 - \theta_2)q_1^S - \text{Area } P - \chi(\boldsymbol{\theta}) \end{aligned}$$

[73] Naturally, it is sufficient that the condition holds only for the range of feasible prices and shocks, that is $\forall\{(p_j,\theta_j) : d_j(p_j,\theta_j) \in [\underline{q}_j, \bar{q}_j]\}$. The effect is strictly positive (negative) if the demand curve is strictly log concave (log convex).

Figure 5: Auctioneer surplus measured with reference to good 2



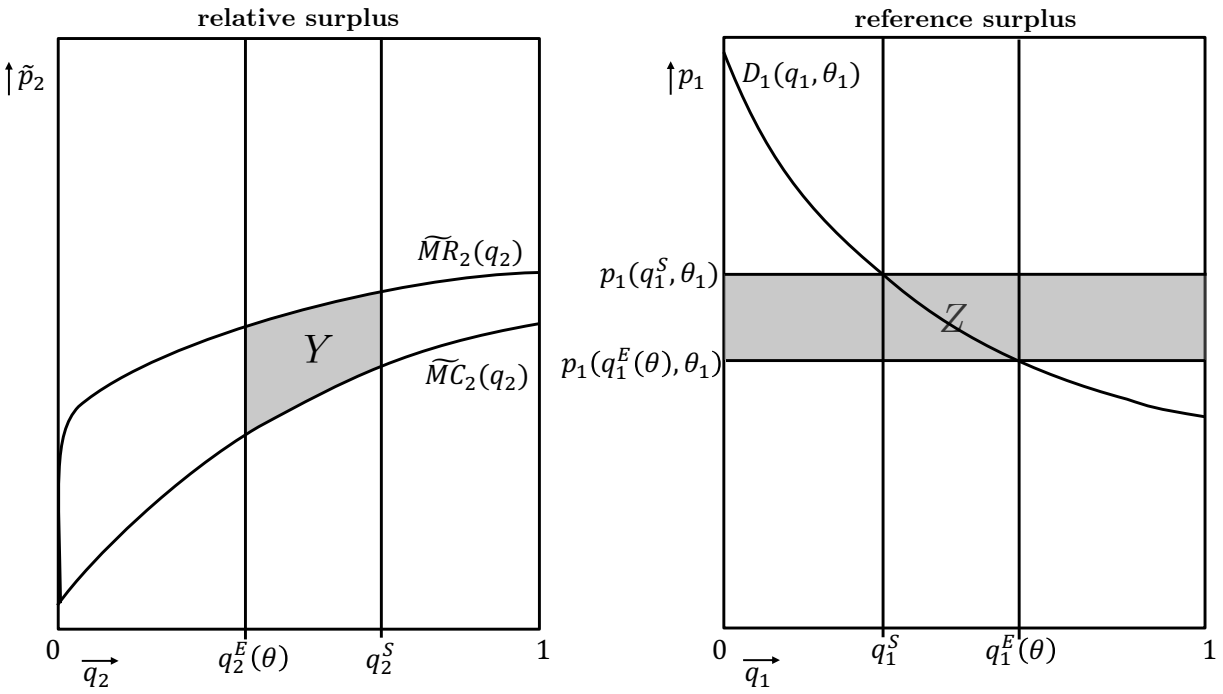
The Auctioneer Relative Price Effect with reference to good 2 is the expected difference in relative surplus between the efficient auction and the SSA, that is, the expected size of Area T. The Auctioneer Price Level Effect with reference to good 2 is the expected difference in reference surplus between the efficient auction and the SSA, that is, the expected size of (Area V + Area W + Area X).

Given shocks $\boldsymbol{\theta}$ for which $q_1^E(\boldsymbol{\theta}) > q_1^S$, the difference in relative surplus between the efficient auction and SSA is Area L+Area U+$(\theta_2-\theta_1)q_1^S$ = Area T+(Area V+Area W+Area X)−$(\theta_1 - \theta_2)q_1^S$ = (Area Q − Area P + Area R − Area S).

D.2 Auctioneer surplus decomposed with reference to good 1

The difference in auctioneer surplus between the efficient auction and SSA can equivalently be measured with reference to good 1. This is illustrated for shocks for which $q_1^E(\boldsymbol{\theta}) > q_1^S$ (and correspondingly $q_2^E(\boldsymbol{\theta}) < q_2^S$) in Figure 6. The difference in relative surplus between the efficient auction and SSA is measured by the adjustment term $(\theta_1 - \theta_2)q_1^S > 0$ less the Area Y. Relative to the SSA, the efficient auction reduces the price level, measured by the price of good 1, reducing the auctioneer's reference surplus by Area Z.

Figure 6: Auctioneer surplus measured with reference to good 1



The Auctioneer Relative Price Effect with reference to good 1 is the expected difference in relative surplus between the efficient auction and the SSA, that is, the expected size of Area Y. The Auctioneer Price Level Effect with reference to good 1 is the expected difference in reference surplus between the efficient auction and the SSA, that is, the expected size of Area Z.

E Example of Proposition 4

Proposition 10 *The Auctioneer Relative Price Effect with reference to good 2 is positive if the two marginal cost curves are independent, that is,* $\frac{\partial^2 C(q_1,q_2)}{\partial q_j \partial q_k} = 0, j,k \in \{1,2\}, j \neq k$*, and the slopes of the inverse marginal cost curves are log concave.*

Proof. Label the inverse marginal cost curves, in the case in which the marginal cost curves are independent, by $mc_j(p_j), j = \{1,2\}$.

Suppose $mc_1'(p_1)$ and $mc_2'(p_2)$ are log concave. We have $q_1 = \widetilde{mc}_1(\widetilde{p}_1)$ and $\widetilde{MC}_1(q_1) = MC_1(q_1) - MC_2(q_2(q_1))$, so

$$MC_1(\widetilde{mc}_1(\widetilde{p}_1)) - MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))) = \widetilde{p}_1$$

Differentiating w.r.t. $\widetilde{p}_1$, noting $MC_j'(q_j) = \frac{1}{mc_j'(MC_j(q_j))}$, $j = \{1,2\}$, and $q_2'(\widetilde{mc}_1(\widetilde{p}_1)) = -1 \; \forall \widetilde{p}_1$,

$$\begin{aligned} &MC_1'(\widetilde{mc}_1(\widetilde{p}_1))\widetilde{mc}_1'(\widetilde{p}_1) - MC_2'(q_2(\widetilde{mc}_1(\widetilde{p}_1)))q_2'(\widetilde{mc}_1(\widetilde{p}_1))\widetilde{mc}_1'(\widetilde{p}_1) = 1 \\ \Longrightarrow &\left(\frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_1'(MC_1(\widetilde{mc}_1(\widetilde{p}_1)))} + \frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_2'(MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))))}\right)\frac{\widetilde{mc}_1'(\widetilde{p}_1)}{\widetilde{mc}_1(\widetilde{p}_1)} = 1 \end{aligned}$$

$\widetilde{mc}_1(\widetilde{p}_1)$ is log concave iff $\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}_1'(\widetilde{p}_1)}{\widetilde{mc}_1(\widetilde{p}_1)}\right) \leq 0$. Differentiating w.r.t. $\widetilde{p}_1$,

$$\begin{aligned} &\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_1'(MC_1(\widetilde{mc}_1(\widetilde{p}_1)))} + \frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_2'(MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))))}\right)\left(\frac{\widetilde{mc}_1'(\widetilde{p}_1)}{\widetilde{mc}_1(\widetilde{p}_1)}\right) \\ &+\left(\frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_1'(MC_1(\widetilde{mc}_1(\widetilde{p}_1)))} + \frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_2'(MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))))}\right)\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}_1'(\widetilde{p}_1)}{\widetilde{mc}_1(\widetilde{p}_1)}\right) = 0 \end{aligned} \tag{4}$$

Noting $\widetilde{mc}_1(\widetilde{p}_1) = q_1$,

$$\begin{aligned} &\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_1'(MC_1(\widetilde{mc}_1(\widetilde{p}_1)))} + \frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc_2'(MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))))}\right) \\ =&\frac{\partial}{\partial q_1}\left(\frac{q_1}{mc_1'(MC_1(q_1))} + \frac{q_1}{mc_2'(MC_2(q_2(q_1)))}\right)\frac{\partial \widetilde{mc}_1(\widetilde{p}_1)}{\partial mc_1} \end{aligned}$$

Define $f_1(q_1) = mc_1'(MC_1(q_1)), f_2(q_1) = mc_2'(MC_2(q_2(q_1)))$ and $g_m(q_1) = \frac{q_1}{f_m(q_1)}$, $h_m(q_1) = f_m(q_1) - q_1 f_m'(q_1), m = \{1,2\}$.

Then $g'_m(q_1) = \frac{1}{(f_m(q_1))^2}\left(f_m(q_1) - q_1 f'_m(q_1)\right)$ and $h'_m(q_1) = -q_1 f''_m(q_1)$. Then $h_m(0) = f_m(0) > 0, m = \{1,2\}$, because the two marginal cost curves are strictly increasing.

$mc'_1(p_1) = f_1(mc_1(p_1))$ and $mc'_2(p_2) = f(q_1(mc_2(p_2)))$ are log concave,

$$\implies \frac{\partial}{\partial p_1}\left(f'_1(mc_1(p_1))\right) = \frac{\partial}{\partial p_1}\left(\frac{mc''_1(p_1)}{mc'_1(p_1)}\right) \leq 0,\ \frac{\partial}{\partial p_2}\left(-f'_2(q_1(mc_2(p_2)))\right) = \frac{\partial}{\partial p_2}\left(\frac{mc''_2(p_2)}{mc'_2(p_2)}\right) \leq 0$$
$$\implies \frac{\partial}{\partial q_1}\left(f'_m(q_1)\right)\frac{\partial mc_m(p_m)}{\partial p_m} = f''_m(q_1)mc'_m(p_m) \leq 0, m = \{1,2\}$$

And $mc_m(p_m), m = \{1,2\}$ are strictly increasing, so $f''_m(q_1) \leq 0, h'_m(q_1) > 0, m = \{1,2\}$. And so,

$$\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc'_1(MC_1(\widetilde{mc}_1(\widetilde{p}_1)))} + \frac{\widetilde{mc}_1(\widetilde{p}_1)}{mc'_2(MC_2(q_2(\widetilde{mc}_1(\widetilde{p}_1))))}\right) \geq 0$$

In addition, inverse relative marginal cost is increasing in $\widetilde{p}_1$, $\widetilde{mc}_1(\widetilde{p}_1) = q_1$ is always non-negative, and the marginal cost curves are strictly increasing, so $\frac{\partial}{\partial \widetilde{p}_1}\left(\frac{\widetilde{mc}'_1(\widetilde{p}_1)}{\widetilde{mc}_1(\widetilde{p}_1)}\right) \leq 0$ follows from Equation 4. The inverse marginal cost curve is log concave, and the rest of the proof follows from Proposition 4. ■

F Generalised Pareto distribution

Suppose bidders' values have a Generalised Pareto distribution. This implies demand curves of the form

$$d_j(p_j, \theta_j) = \left(1 + \frac{\xi_j(p_j - \theta_j - \mu_j)}{\sigma_j}\right)^{-\frac{1}{\xi_j}},\ j = \{1,2\}$$

where $\xi_j < 1$, $\sigma_j > 0$.

The expected price elasticity of demand for good $j \in \{1,2\}$ is

$$\mathbb{E}\left[\epsilon_j^D(q_j, \theta_j)\right] = \frac{D_j(q_j, \theta_j) - \theta_j}{\frac{\partial D_j(q_j, \theta_j)}{\partial q_j} q_j} = \left(\frac{1}{\xi_j} - \frac{\mu_j}{\sigma_j}\right) q_j^{\xi_j} - \frac{1}{\xi_j}$$

and

$$\frac{\partial \mathbb{E}\left[\epsilon_j^D(q_j, \theta_j)\right]}{\partial q_j} = \xi_j\left(\frac{1}{\xi_j} - \frac{\mu_j}{\sigma_j}\right) q_j^{\xi_j - 1}$$

The expected price elasticities of demand are negative. If $\{\xi_1, \xi_2\} < 0$, this implies $\frac{1}{\xi_j} - \frac{\mu_j}{\sigma_j} \leq 0 \ \forall j = \{1, 2\}$. Under this condition, the expected price elasticities are therefore both increasing in the quantity sold. And so

$$\mathbb{E}\left[\epsilon_1^D(0, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1, \theta_2)\right] \geq 0 \Longrightarrow \mathbb{E}\left[\epsilon_1^D(q_1, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1 - q_1, \theta_2)\right] \geq 0 \forall q_1 \in [0, 1]$$

We have $\mathbb{E}\left[\epsilon_1^D(0, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1, \theta_2)\right] = -\frac{1}{\xi_1} + \frac{\mu_2}{\sigma_2}$. So, if $\{\xi_1, \xi_2\} < 0$ and $\frac{\mu_2}{\sigma_2} \geq \frac{1}{\xi_1}$, then $\mathbb{E}\left[\epsilon_1^D(q_1, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1 - q_1, \theta_2)\right] \geq 0 \forall q_1 \in [0, 1]$.

Now turn to the case of $\{\xi_1, \xi_2\} > 0$. If either (i) the expected price elasticities are both increasing in the quantity sold and $\mathbb{E}\left[\epsilon_1^D(1, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(0, \theta_2)\right] \leq 0$ or (ii) the expected price elasticities are both decreasing in the quantity sold and $\mathbb{E}\left[\epsilon_1^D(0, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1, \theta_2)\right] \leq 0$, then $\mathbb{E}\left[\epsilon_1^D(q_1, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1 - q_1, \theta_2)\right] \leq 0 \forall q_1 \in [0, 1]$. We have

$$\begin{aligned}
\mathbb{E}\left[\epsilon_1^D(1, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(0, \theta_2)\right] &= \frac{1}{\xi_2} - \frac{\mu_1}{\sigma_1} \\
\mathbb{E}\left[\epsilon_1^D(0, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1, \theta_2)\right] &= -\frac{1}{\xi_1} + \frac{\mu_2}{\sigma_2} \\
\frac{\partial \mathbb{E}\left[\epsilon_j^D(q_j, \theta_j)\right]}{\partial q_j} &\begin{cases} \geq 0 \forall q_j & \text{if } \left(\frac{1}{\xi_j} - \frac{\mu_j}{\sigma_j}\right) \geq 0 \\ \leq 0 \forall q_j & \text{if } \left(\frac{1}{\xi_j} - \frac{\mu_j}{\sigma_j}\right) \leq 0 \end{cases}
\end{aligned}$$

for $j \in \{1, 2\}$. And so, $\mathbb{E}\left[\epsilon_1^D(0, \theta_1)\right] - \mathbb{E}\left[\epsilon_2^D(1, \theta_2)\right] \leq 0$ if (i) $\frac{1}{\xi_1} \geq \frac{\mu_1}{\sigma_1}, \frac{1}{\xi_2} \geq \frac{\mu_2}{\sigma_2}$ and $\frac{\mu_1}{\sigma_1} \geq \frac{1}{\xi_2}$ or (ii) $\frac{1}{\xi_1} \leq \frac{\mu_1}{\sigma_1}, \frac{1}{\xi_2} \leq \frac{\mu_2}{\sigma_2}$, and $\frac{\mu_2}{\sigma_2} \leq \frac{1}{\xi_1}$. Or equivalently, either

$$\text{(i) } \frac{\mu_2}{\sigma_2} \leq \frac{1}{\xi_2} \leq \frac{\mu_1}{\sigma_1} \leq \frac{1}{\xi_1}; \text{ or } \quad \text{(ii) } \frac{1}{\xi_2} \leq \frac{\mu_2}{\sigma_2} \leq \frac{1}{\xi_1} \leq \frac{\mu_1}{\sigma_1}$$

G Mexican government debt auctions: alternative estimation method

In the main analysis, I estimate the distribution of residualised bids (equal to the difference between bids and the secondary market price) that the government expects to face in auction t by sampling with replacement from the set of bids that bidders submit in auction $t-1$. By pooling bids from just one auction, this approach is consistent even if there is unobserved heterogeneity in the distribution of bids across auctions. Nonetheless, by resampling from the set of individual bidders, I assume that their valuations are identically distributed.[74]

In this appendix, I use an alternative resampling procedure to estimate the differences in surpluses between the counterfactual PMA and the actual mechanism. The approach relaxes the assumption that bidders' valuations are identically distributed, but assumes that the distribution of residualised bids is stable over time.

The set of $T-1$ auctions in the sample, which excludes auction t, represent a set of $T-1$ realisations of the aggregate bid functions for bonds $j = \{1, 2\}$. I simulate one possible realisation of residualised bids by drawing the pair of aggregate bid functions for bonds $j = \{1, 2\}$ from this sample, and subtracting the secondary market prices on the day of the corresponding auction. Repeating this procedure a large number of times simulates the distribution of residualised bids the government expects to face. The method is otherwise identical to the main analysis.

The results using this method for the cases in which the objectives are to maximise welfare and auctioneer surplus, respectively, are shown in Tables 6 and 7.

[74] A third alternative approach would be to group bidders into a finite number of groups and resample within group.

Table 6: Surplus estimates and model predictions in Mexican CETES primary auctions given an objective to maximise welfare

Pair of bonds		28- & 91-day	182- & 364-day
Bond 1	(term, days)	28	364
Bond 2	(term, days)	91	182
Issuance of bond 1 in SSA	(mean, % of total)	39.533	53.775
		(0.405)	(0.914)
	(std. dev., % of total)	5.177	7.824
		(0.323)	(0.576)
Bidder surplus in SSA	(bps of issuance)	0.767	2.709
		(0.037)	(0.205)
	(MEX$ million per auction)	1.248	6.362
		(0.058)	(0.512)
Difference between counterfactual PMA and actual SSA			
Δ Issuance of bond 1	(mean, % of total)	1.735	0.121
		(2.064)	(3.453)
	(std. dev., % of total)	26.599***	29.798***
		(1.099)	(1.789)
Δ Welfare	(bps of issuance)	0.220***	1.027***
		(0.031)	(0.141)
	(MEX$ million per auction)	0.357***	2.411***
		(0.051)	(0.330)
Δ Bidder surplus	(bps of issuance)	0.063*	0.855***
		(0.026)	(0.248)
	(MEX$ million per auction)	0.102*	2.008***
		(0.043)	(0.580)
Δ Auctioneer surplus	(bps of issuance)	0.157***	0.172
		(0.033)	(0.231)
	(MEX$ million per auction)	0.255***	0.403
		(0.054)	(0.543)
Bidder Uncertainty Effect	(bps of issuance)	0.081**	0.923***
		(0.030)	(0.240)
	(MEX$ million per auction)	0.132**	2.168***
		(0.049)	(0.562)
Bidder Allocation Effect	(bps of issuance)	-0.018	-0.068
		(0.013)	(0.039)
	(MEX$ million per auction)	-0.030	-0.160
		(0.021)	(0.093)

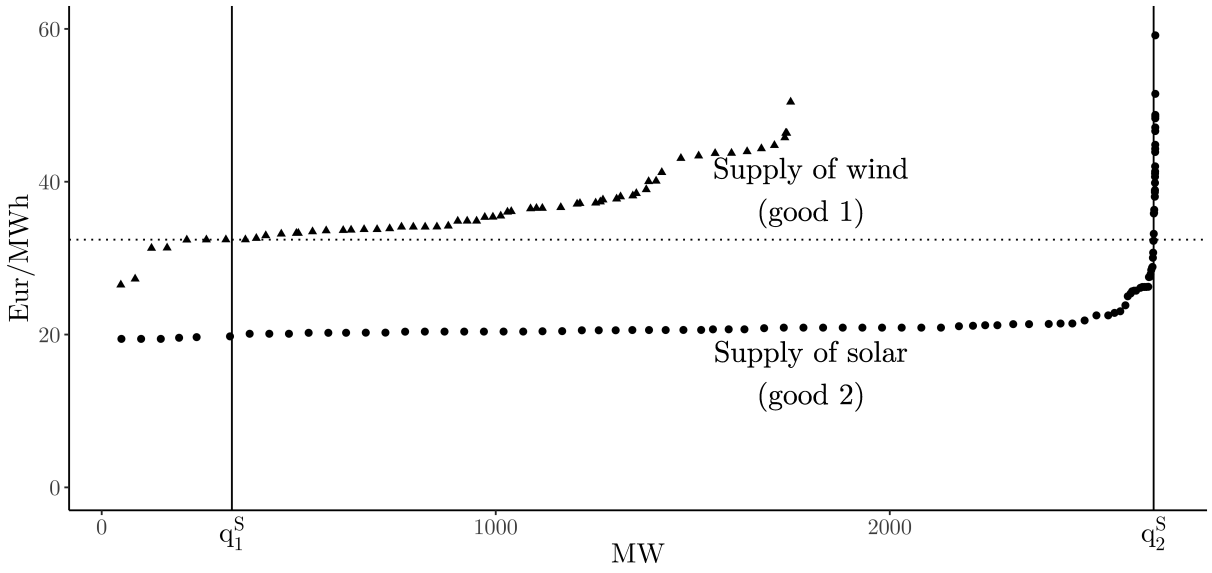
Note: bps denote basis points. Surpluses measured in bps are weighted by total quantities issued and surpluses measured in MEX$ are unweighted. Standard errors are in parentheses. The level of welfare and auctioneer surplus in the actual SSA, which depend on the auctioneer's unobserved total benefit from issuing the total supply, are unidentified. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 7: Surplus estimates and model predictions in Mexican CETES primary auctions given an objective to maximise auctioneer surplus

Pair of bonds		28- & 91-day	182- & 364-day
Bond 1	(term, days)	91	364
Bond 2	(term, days)	28	182
Issuance of bond 1 in SSA	(mean, % of total)	39.585	53.893
		(0.427)	(0.893)
	(std. dev., % of total)	5.483	7.643
		(0.320)	(0.524)
Bidder surplus in SSA	(bps of issuance)	0.775	2.674
		(0.039)	(0.211)
	(MEX$ million per auction)	1.261	6.280
		(0.061)	(0.527)
Difference between counterfactual PMA and actual SSA			
Δ Issuance of bond 1	(mean, % of total)	7.067***	2.115
		(1.875)	(2.464)
	(std. dev., % of total)	23.884***	21.463***
		(1.153)	(1.558)
Δ Welfare	(bps of issuance)	0.193***	0.712***
		(0.029)	(0.140)
	(MEX$ million per auction)	0.314***	1.672***
		(0.048)	(0.329)
Δ Bidder surplus	(bps of issuance)	-0.188***	-0.461**
		(0.023)	(0.150)
	(MEX$ million per auction)	-0.306***	-1.084**
		(0.036)	(0.353)
Δ Auctioneer surplus	(bps of issuance)	0.382***	1.173***
		(0.036)	(0.150)
	(MEX$ million per auction)	0.621***	2.756***
		(0.058)	(0.355)

Note: bps denote basis points. Surpluses measured in bps are weighted by total quantities issued and surpluses measured in MEX$ are unweighted. Standard errors are in parentheses. The level of welfare and auctioneer surplus in the actual SSA, which depend on the auctioneer's unobserved total benefit from issuing the total supply, are unidentified. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 7: Aggregate supply curves for wind and solar energy



Supply curves of wind and solar energy, estimated by Fabra and Montero (2023) using data on potential investment projects in Spain with applications for planning permission at the Registry of Renewables Installations in Spain (RIPRE) in January 2019 until March 2020.

H Spanish renewable energy procurement: model predictions

The large number of current and planned auctions for renewable energy contracts across a range of countries (IRENA, 2013) exhibit large variation in design, with a broad trend towards multi-product auctions of multiple technologies (Szabó et al., 2020). In these energy procurement auctions, the auctioneer acts as the buyer and the bidders act as sellers. This is a mirror image of the main model, and the results are of course unchanged. In many cases, the auctions are special cases of PMAs, either with the two technologies treated interchangeably, or with a demand curve expressing the government's preference for diversified procurement.

I illustrate the model's predictions for Fabra and Montero's (2023) analysis of Spanish renewable energy auctions held in January 2021, in which solar and wind producers were permitted to bid for energy contracts within a single auction. The government's total demand was fixed, and in the neighbourhood of the observed quantities procured of the two goods, the two types of energy were treated interchangeably, so the auctioneer's revealed preference was perfect substitutes.[75]

[75]The auction rules specified minimum quotas reserved for each energy type, but the auction outcome

Using data on 2019 applications for renewable energy projects, Fabra and Montero (2023) estimate the aggregate bid functions, i.e., the supply curves, for solar and wind energy in Spain, which I use to determine the properties of the supply curves relevant to the model predictions. Figure 7 shows these bid functions in the case of zero shocks.

While the supply curve for wind is approximately linear, the inverse supply curve for solar is convex in the neighbourhood of the observed pair of quantities procured, so wind is labelled good 1 in the model by Lemma 3 in Appendix B. Given the supply estimates, the PMA run by the Spanish government therefore increases the expected quantity procured of wind energy relative to the pair of separate simultaneous auctions that would maximise expected welfare.

Both the Auctioneer Price Level Effect and Auctioneer Relative Price Effect with reference to good 1 (wind) equal zero, given the linearity of the supply curve for wind and the auctioneer's perfect substitute preferences. So the model predicts that the government is indifferent between the PMA and running separate auctions.

The supplies of both goods are log concave, so that bidders prefer the uncertainty in the quantities procured of the two goods in the PMA, and, by Proposition 1, the Bidder Uncertainty Effect is positive.[76]

Because of the auctioneer's indifference, the positive Bidder Uncertainty Effect must dominate so that the PMA increases expected bidder surplus, as it increases expected welfare.

In this setting, the model predicts that the entirety of the welfare gain of the PMA translates into higher expected bidder surplus.

triggered neither quota. Note the minimum quota design can be viewed as a PMA in which the demand curve is a step function.

[76]The Bidder Allocation Effect is ambiguous. At the quantities procured in the SSA in Fabra and Montero's (2023) simulation (denoted by q_1^S and q_2^S in Figure 7, which are the pair that maximises expected surplus at $\theta_1 - \theta_2 = \mathbb{E}[\theta_1 - \theta_2]$) the supply of wind energy is more price elastic than the supply of solar and the auctioneer's relative demand for good 1 is zero at the quantities procured in the SSA (given the auctioneer views the two goods as perfect substitutes). Proposition 2 therefore does not apply as the supply curves are log concave.

Conclusion

This thesis focuses on two themes relating to auctions in financial markets. The first two chapters analyse the the behaviour of participants and assess the ability of alternative models to explain bidding data from the Bank of England's liquidity auctions. The final two chapters analyse the effects on welfare, and the distribution of this surplus between bidders and the auctioneer, of alternative auction designs. In these chapters, I compare designs in the contexts of the Bank of England's liquidity auctions and Mexico's domestic government bond market, both theoretically and empirically.

The first two chapters contribute to research applying structural approaches to questions relating to multi-unit auctions. This literature is motivated by the fact that theoretical predictions are often ambiguous for multi-unit auctions. For example, while the theoretical literature provides clear predictions of the revenue and efficiency rankings of various auction rules for single-unit auctions, fewer analogous results exist for multi-unit auctions. Neither Vickrey's (1961) revenue equivalence theorem nor Milgrom and Weber's (1982) revenue ranking results hold in multi-unit auctions with multi-unit demands (Ausubel et al., 2014). For example, there is no definitive theoretical ranking in terms of efficiency or revenue between uniform-price and pay-as-bid auctions. Chapters 1 and 2 contribute to the growing empirical literature by building on existing structural models and developing new identifying conditions to explain bidding behaviour in the context of the Bank of England's liquidity auctions.

In the final two chapters, I focus on the benefits of alternative auction designs when the auctioneer wishes to sell multiple substitute goods. While selling goods using multi-product multi-unit auctions can increase efficiency, the effect on the distribution of surplus between bidders and the auctioneer is less clear. Chapters 3 and 4 contribute to a relatively new area of research and my work is closest to that of Fabra and Montero (2023). I provide new conditions under which increasing efficiency by combining competitive auctions of substitute goods increases both expected bidder and auctioneer surplus relative to running separate simultaneous auctions, as well as new conditions under which one side's expected surplus falls. I also contribute to this literature empirically, by estimating the magnitude of the effects in the Bank of England's liquidity auctions and Mexico's domestic government debt market.

References

Aguirre, I., Cowan, S., and Vickers, J. (2010). Monopoly Price Discrimination and Demand Curvature. *American Economic Review, 100*(4), 601-615.

Allen, J., Kastl, J., and Wittwer, M. (2022). Interconnected multi-unit auctions: An empirical analysis. Working Paper.

Alphandary, A. (2014). Risk managing loan collateral at the Bank of England. *Bank of England Quarterly Bulletin, 54*(2), 190–201.

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

Armantier, O., Holt, C. A., and Plott, C. R. (2013). A procurement auction for toxic assets with asymmetric information. *American Economic Journal: Microeconomics, 5*(4), 142–162.

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

Ausubel, L., Cramton, P., Feliz-Ozbay, E., Higgins, N., Ozbay, E., and Stocking, A. (2013). Common-Value Auctions with Liquidity Needs: An Experimental Test of a Troubled Assets Reverse Auction. In *Handbook of Market Design*. Oxford University Press.

Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., and Weretka, M. (2014). Demand Reduction and Inefficiency in Multi-Unit Auctions. *Review of Economic Studies, 81*(4), 1366-1400.

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

Balibek, E., and Memis, H. A. (2012). Turkish Treasury Simulation Model for Debt Strategy Analysis. *World Bank Policy Research Working Paper 6091.*

Banerjee, R., and Mio, H. (2018). The impact of liquidity regulation on banks. *Journal of Financial Intermediation, 35*, 30-44.

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.

Bank of England (2022). The Bank of England's Asset Purchase Facility: APF Operating Procedures; Corporate Bond Purchase Scheme; Gilt Purchases. `https://www.bankofengland.co.uk/-/media/boe/files/markets/asset-purchase-facility/operating-procedures.pdf`.

Bergstrom, P., Holmlund, A., and Lindberg, S. (2002). The SNDO's Simulation Model for Government Debt Analysis. Swedish National Debt Office mimeo.

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

Bolder, D. J., and Deeley, S. (2011). The Canadian Debt Strategy Model: An Overview of the Principal Elements. Bank of Canada Discussion Paper.

Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market. *American Economic Review, 92*(5), 1376–1405.

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

Bulow, J., and Klemperer, P. (2012). Regulated Prices, Rent Seeking, and Consumer Surplus. *Journal of Political Economy, 120*(1), 160-186.

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

Castellanos, S., and Oviedo, M. (2008). Optimal Bidding in the Mexican Treasury Securities Primary Auctions: Results of a Structural Econometric Approach. *Latin American Journal of Economics* (formerly *Cuadernos de Economía*), *45*(131), 3-28.

Chakraborty, I., and Engelbrecht-Wiggans, R. (2014). Truthful Bidding in Uniform-Price Multi-Unit Auctions. Working paper.

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

Chen, D., and Duffie, D. (2021). Market fragmentation. *American Economic Review, 111*(7), 2247-2274.

Chen, S., Ganum, P., Liu, L., Martinez, L., and Peria, S. M. (2019). Debt Maturity and the Use of Short-Term Debt: Evidence from Sovereigns and Firms. IMF Departmental Papers / Policy Papers, No 2019/001, International Monetary Fund.

Cole, H. L., Neuhann, D., and Ordoñez, G. (2022). Asymmetric Information and Sovereign Debt: Theory Meets Mexican Data. *Journal of Political Economy, 130*(8), 2055-2109.

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

Duarte, M., Magnolfi, L., and Roncoroni, C. (2021). The competitive conduct of consumer cooperatives. Working paper.

Fabra, N., and Montero, J. P. (2023). Technology-Neutral Versus Technology-Specific Procurement. *The Economic Journal, 133*(650), 669–705.

Fabinger, M., and Weyl, E. G. (2012). Pass-through and demand forms. Working paper.

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* [Speech]. Bank of England.

Frost, T., Govier, N., and Horn, T. (2015). Innovations In The Bank's Provision Of Liquidity Insurance Via Indexed Long-Term Repo (ILTR) Operations. *Bank Of England Quarterly Bulletin, 55*(2), 181–188.

Gobierno de México (2022). Plan Anual de Financiamiento 2022, Secretaría De Hacienda y Crédito Público. `https://www.finanzaspublicas.hacienda.gob.mx/work/models/Finanzas_Publicas/docs/paquete_economico/paf/paf_2022.pdf`.

Gentry, M. L., Hubbard, T. P., Nekipelov, D., and Paarsch, H. J. (2018). Structural

Econometrics of Auctions: A Review. *Foundations and Trends in Econometrics, 9*(2-4), 79–302.

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.

Greenwood, R., Hanson, S. G., and Stein, J. C. (2015). A Comparative-Advantage Approach to Government Debt Maturity. *Journal of Finance, 70*(4), 1683-1722.

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

Gupta, S., and Lamba, R. (2022). Treasury auctions during a crisis. Working paper.

Haile, P. (2001). Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales. *American Economic Review, 91*(3), 399–427.

Haile, P., Hong, H., and Shum, M. (2003). Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions. NBER Working Paper 10105.

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

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

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

Hortaçsu, A., and McAdams, D. (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 McAdams, D. (2018). Empirical Work on Auctions of Multiple Objects. *Journal of Economic Literature, 56*(1), 157-184.

Hortaçsu, A., and Perrigne, I. (2021). Empirical perspectives on auctions. In *Handbook of*

Industrial Organization, 5(1) 81-175. Elsevier.

Hortaçsu, A., and Puller, S. L. (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.

IRENA (2013). Renewable Energy Auctions in Developing Countries. `www.irena.org/publications/2013/Jun/Renewable-Energy-Auctions-in-Developing-Countries`.

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), 339–352.

Klemperer, P. (2008). A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions. Mimeo: 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), 526-536.

Klemperer, P. (2018). Product-Mix Auctions. Working Paper 2018-W07, Nuffield College.

Klemperer, P., and Meyer, M. (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica, 57*(6), 1243–1277.

Lings, S. (2013). Welfare Effects of Combining Auctions. M.Phil Thesis: University of Oxford.

McAdams, D. (2008). Partial Identification and Testable Restrictions in Multi-Unit Auctions. *Journal of Econometrics, 14*6(1), 74–85.

Milgrom, P., and Weber, R. J. (1982). A Theory of Auctions and Competitive Bidding. *Econometrica, 50*(5), 1089-1122.

Milnes, A. (2010). Creating Confidence in Cash. *Blueprint*.

Missale, A., and Blanchard, O. J. (1994). The debt burden and debt maturity. *American Economic Review, 84*(1), 309-319.

Nixon, D. (2014). Sterling Monetary Framework Annual Report 2013–14. *Bank of England*

Quarterly Bulletin, 54(2), 218-225.

Pick, A., and Anthony, M. (2006). A simulation model for the analysis of the UK's sovereign debt strategy. UK Debt Management Office Research Paper.

Rasmussen, C. E., and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press. Cambridge, Massachusetts.

Reguant, M. (2014). Complementary bidding mechanisms and startup costs in electricity markets. *Review of Economic Studies, 81*(4), 1708-1742.

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

Rostek, M. J., and Yoon, J. H. (2021). Innovation in Decentralized Markets: Synthetic Products vs. Trading Technology. Working Paper.

Roussille, N., and Scuderi, B. (2021). Bidding for Talent: Equilibrium Wage Dispersion on a High-Wage Online Job Board. Working paper.

Song, Z., and Zhu, H. (2018). Quantitative easing auctions of treasury bonds. *Journal of Financial Economics, 128*(1), 103-124.

Starc, A., and Wollmann, T. G. (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.

Szabó, L., Bartek-Lesi, M., Dézsi, B., Diallo, A., and Mezősi, A. (2020). Auctions for the support of renewable energy: Lessons learnt from international experiences – Synthesis report of the AURES II case studies. *D2.3. Aures II Project* `http://aures2project.eu/`.

Umlauf, S. R. (1993). An empirical study of the Mexican Treasury bill auction. *Journal of Financial Economics, 33*(3), 313-340.

Vasserman, S., and Watt, M. (2021). Risk aversion and auction design: Theoretical and empirical evidence. *International Journal of Industrial Organization, 79*, 102758.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal*

of Finance, 16(1), 8–37.

Waugh, F. V. (1944). Does the consumer benefit from price instability?. *Quarterly Journal of Economics, 58*(4), 602-614.

Weitzman, M. (1974). Prices vs. Quantities. *Review of Economic Studies, 41*(4), 477–491.

Weyl, E. G., and Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy, 121*(3), 528-583.

Wilson, R. (1979). Auctions of Shares. *Quarterly Journal of Economics, 93*(4), 675–689.

Winters, B. (2012). Review of the Bank of England's Framework for Providing Liquidity to the Banking System. Bank of England.

Wittwer, M. (2021). Connecting disconnected financial markets?. *American Economic Journal: Microeconomics, 13*(1), 252-282.

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.