A New Auction for Substitutes:
Central-Bank Liquidity Auctions, "Toxic Asset" Auctions, and Variable Product-Mix Auctions

Paul Klemperer, 2008, revised 2009

Nuffield College, Oxford, OX1 1NF, England
paul.klemperer@economics.ox.ac.uk
+44 777 623 0123

DRAFT –COMMENTS WELCOME

The most recent public version of this paper is available at http://www.paulklemperer.org

Abstract

I describe a new static (sealed-bid) auction for multiple substitute goods. As in a two-sided simultaneous multiple round auction (SMRA), bidders bid on multiple assets simultaneously, and bid-takers choose supply functions across assets. The auction yields more efficiency, revenue, information, and trade than running multiple separate auctions, but is often simpler to use and understand, and less vulnerable to collusion, than a SMRA. I designed it after the 2007 Northern Rock bank run to help the Bank of England fight the credit crunch; in 2008 the U.S. Treasury planned (but later cancelled) using a related design to buy "toxic assets".

Keywords: multi-object auction, TARP, central banking, simultaneous ascending auction, treasury auction, term auction, toxic assets.

JEL Nos. D44 (Auctions), E58 (Central Banking)

I have advised the Bank of England and the U.S. Treasury and have been consulted by other Central Banks, government agencies, etc., about these issues. I thank the relevant officials for help, but the views here are my own and do not represent those of any organisation. I am particularly grateful to Jeremy Bulow and Daniel Marszalec for their help in advising the Bank of England, and I also especially thank Eric Budish, Vince Crawford, Meg Meyer, Marco Pagnozzi, Rakesh Vohra and many other friends and colleagues for helpful advice .... [Acknowledgements to be completed later]

1. Introduction

How should goods that both seller(s) and buyers view as imperfect substitutes be sold, especially when multi-round auctions are impractical?

This was the Bank of England's problem in autumn 2007 as the credit crunch began.[1] The Bank urgently wanted to supply liquidity to banks, and was therefore willing to accept a wider-than-usual range of collateral, but it wanted a correspondingly higher interest rate against any weaker collateral it took. Furthermore, because financial markets move fast, any auction had to take place at a single instant. In a multi-stage auction bidders who had entered the highest bids early on might change their minds about wanting to be winners before the auction closed,[2] and the financial markets might themselves be influenced by the evolution of the auction, which magnifies the difficulties of bidding and invites manipulation.

An equivalent problem to the Bank of England's is that of a firm which can supply multiple varieties of a product (at different costs), but with a total capacity constraint, to customers with different strengths of preferences between those product varieties, and where transaction costs or other time pressures make multiple-round auctions infeasible.[3] (The different varieties of a product could include different points of delivery, different warranties, or different restrictive covenants on use.)

A similar problem was the U.S. Treasury's autumn 2008 Troubled Asset Recovery Program (TARP) plan to spend up to $700 billion buying "toxic assets" with a face value well in excess of $1 trillion. There were of the order of 25,000 closely-related but distinct

[1] The crisis began in early August 2007, and a bank run led to Northern Rock's collapse in mid-September. Immediately subsequently, the Bank of England ran four auctions to supply additional liquidity to banks -- but received no bids in any of them (for reasons that are not the subject of this note). The Bank then consulted me, and I got assistance from Jeremy Bulow and Daniel Marszalec. Starting later in 2007, the Bank ran additional simple (more successful) auctions while consulting on the ideas discussed below.
[2] Some evidence is that most bids in standard Treasury auctions are made in the last few minutes, and a large fraction in the last few seconds. For a multi-round auction to have any merit, untopped bids cannot be withdrawn without incurring penalties.
[3] That is, the Bank can be thought of as a "firm" whose "product" is loans, which come in different "varieties" corresponding to the different collaterals they are made against, and the total supply of which may be constrained. The Bank's "customers" are its counterparties, and the "prices" they bid are interest rates.

assets (different "Alt-A" and sub-prime non-agency mortgage-backed securities originally rated AAA), and perhaps 300 likely sellers, but the largest 10 held of the order of two-thirds of the total volume. And, as above, the volatility of financial markets, their sensitivity to news, and the possibility of feedback effects between them and any multi-round auction made a dynamic auction problematic.

This paper outlines a solution to all these problems. I first developed it for the Bank of England,[4] and later made a similar proposal to the U.S. Treasury (which would probably have used a related design if it had not abandoned its plans to buy toxic assets).[5]

My design is straightforward in concept: allow each bidder to offer one or more *sets* of bids; each set contains one or more bids; each bid specifies a price (or, in the Bank of England's auction, an interest-rate) for a quantity of a specific "variety"; and the bids in each set are mutually exclusive. The auctioneer looks at all the sets of bids and then chooses its preferred (uniform) prices, one for each "variety". From each set of bids offered by each bidder, the auctioneer accepts the one that gives the bidder the greatest surplus evaluated at those prices, or no bid if all the bids would give the bidder negative surplus.

The idea is that the menu of mutually-exclusive bids allows each bidder to approximate a demand function, so bidders can, in effect, decide how much of each variety to buy *after*

[4] See note 1. The Bank has held consultations on the proposal, and used some simple versions, but the continuing unsettled state of the financial markets has meant repeated delays. I do *not* give full details of the Bank of England's objectives, constraints, and concerns here, and not all the issues I discuss are relevant to that Bank -- the solution I describe contains more features than needed for most Central Banking applications.

[5] Subsequent to proposing my solution to the Bank of England, I learned that Paul Milgrom was independently pursuing related ideas. He and I therefore made a joint proposal to the U.S. Treasury, together with Jeremy Bulow and Jon Levin in September-October 2008. Other consultants, too, proposed a static (sealed-bid) design, and the Treasury was planning to run a first set of sealed-bid auctions, each for a related group of assets, when it suddenly abandoned its plans to buy subprime assets (in November 2008). Note, however, Larry Ausubel and Peter Cramton (who played an important role in demonstrating the value and viability of using auctions for TARP, see, e.g., Ausubel and Cramton (2008), and ran experiments, see Ausubel et al. (2008)) were also influential advisors, and their proposal to run some dynamic auctions at a later stage was still being explored.

Milgrom (forthcoming) shows how to represent a wide range of bidders' preferences very elegantly while at the same time restricting to substitutable preferences. He also shows his highly-efficient linear-programming approach yields an integer allocation when demands and constraints are integer -- this could be important in some applications, although not in a context such as that of the Bank of England, for which my proposal seems more straightforward and transparent.

seeing the prices chosen. Meanwhile the auctioneer can look at demand *before* choosing the prices; allowing it to choose the prices ex-post creates no problem here, because it meets the demand that each bidder would have chosen for itself given those prices.[6] [7] Also importantly, offers for each variety provide a competitive discipline on the offers for the other varieties, because they are all being auctioned simultaneously.

Compare this with the "standard" approach of running a separate auction for each different "variety".[8] In this case, outcomes are erratic and inefficient, because the auctioneer has to choose how much of each variety to offer before learning bidders' preferences, while bidders have to guess how much to bid for in each auction without knowing what the price-differences between varieties will turn out to be; the wrong bidders may win, and those who do win may be inefficiently allocated across varieties. Furthermore, each individual auction is much more sensitive to market power, to manipulation, and to informational asymmetries, than if all offers competed directly with each other in a single auction. The auctioneer's revenues are correspondingly generally lower. Thus, for example, if the U.S. Treasury had simply predetermined the amount of each type of security to purchase, ignoring the information about demand for the large number of closely-related securities, competition would have been inadequate because of the highly-concentrated ownership of the assets. All these problems also reduce the auctions' value as a source of information. They may also reduce participation, which can create "second-round" feedback effects furthering magnifying the problems.[9]

---

[6] That is, it chooses prices like a Walrasian auctioneer who is equating bidders' demand with the bid-taker's supply in a decentralized process (in which the privately-held information needed to determine the allocation is directly revealed by the choices of those who hold it).

[7] This assumes the conditions for "sincere bidding" are satisfied -- see below.

[8] The Bank of England used this approach (specifying a single collateral quality for each separate auction) pending development of our proposal.

An alternative that can work well in some contexts is to set fixed price supplements for "superior" varieties, and then auction all units as if they are homogenous. But the auctioneer cannot then express any preferences about the proportions of different varieties transacted. (This might be important if, for example, the different varieties are deliveries at locations that are close to different of the firm's plants.) And a Central Bank might not want to signal its view of appropriate price-differentials for different collaterals, even if it felt sufficiently informed to have one. Also, if the U.S. Treasury had simply developed a "reference price" for each asset, and bought those assets that were offered cheapest relative to their reference prices, it would have purchased large quantities of any assets whose reference prices were set too high -- and mistakes would have been inevitable, since the government had so much less information than the sellers.

[9] The feedback effects by which low participation reduces liquidity, which further reduces participation and liquidity, etc., are of much greater concern when there are multiple agents on both sides of the market -- see Klemperer (2008).

The question, of course, is whether our approach can actually be implemented, and -- crucially -- whether it can be done in a way that is simple and robust, and easy for bidders to understand, so that they are happy to participate. We now show this is feasible.

Section 2 provides a simple graphical illustration of our solution for the two-good case. Section 3 discusses extensions. In particular, it is easy to include multiple buyers and multiple sellers, and "swappers" who may be on either, or both, sides of the market. Section 4 observes that our design is essentially a two-sided "proxy" implementation of a standard simultaneous multiple-round auction (SMRA),[10] and contrasts our approach with alternatives. Section 5 concludes.

2. A Simple Two-Variety Example

The application this auction was originally designed for provides a simple illustration. A single seller, the Bank of England (henceforth "the Bank") auctions just two "goods", namely a loan of funds secured against strong collateral, and a loan of funds secured against weak collateral. For simplicity we refer to the two goods as "strong" and "weak".[11]

Each bidder can make any number of bids.[12] Each bid specifies a per-unit price (i.e., interest-rate) for *each* variety and a *single* quantity. Thus, for example, one bidder might make three separate bids: a bid for £375 million at {5.95% for (funds secured against) weak OR 5.70% for (funds secured against) strong}, a bid for an additional £500 million at {5.75% for weak collateral OR 5.5% for strong}, and a bid for a further £300 million at {5.7% for weak OR 0% for strong}. The auctioneer will accept at most one of the two offers from any bid, namely the one that gives the bidder the greatest surplus as measured

[10] See Milgrom (2000) for a description of the SMRA. SMRAs are commonly used for auctioning radio spectrum see, for example, Binmore and Klemperer (2000). See Klemperer (2004) for an introduction to the literature on auctions.
[11] "Strong" corresponds to the "OMO" or "ordinary" collateral the Bank of England traditionally accepted in its "open market operations". "Weak" corresponds to the "wider" or "extended" collateral that the Bank was willing to lend against in the stressed circumstances that developed from autumn 2007.
[12] About 40 counterparties (commercial banks, building societies, etc.) could bid in the Bank's auctions.

by the difference between the price the bidder offers and the market-clearing price, so a "truthful" bidder knows that its preferred offer -- if any -- will be accepted from each of its bids. Note that since offers at a price of zero will never be selected, the last of the three bids above is equivalent to a traditional bid on only a single collateral (weak).[13]

Figure 1a. A Possible Allocation of Funds

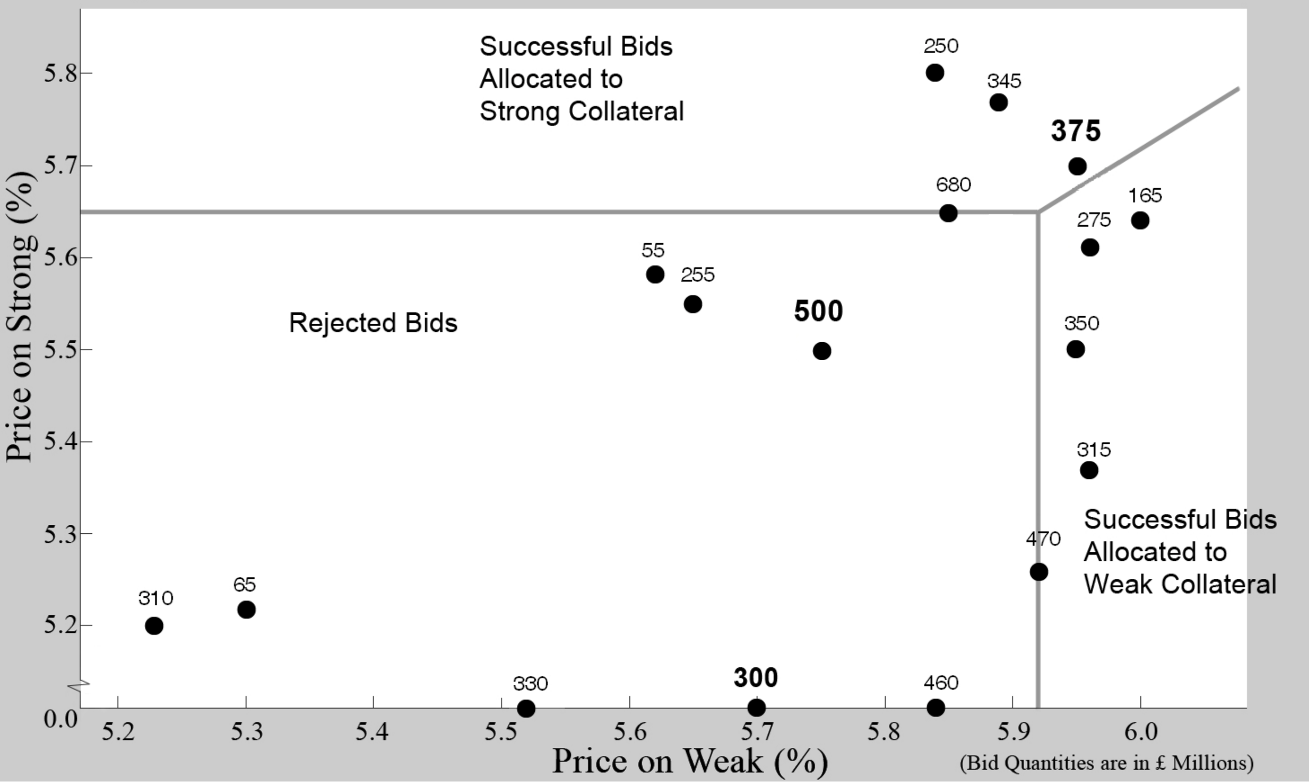

An example of the universe of all the bids submitted by all the bidders is illustrated in Figure 1a. The prices (i.e., interest rates) for weak and strong are plotted vertically and horizontally, respectively, so that each dot in the chart represents an "either/or" bid.[14] The number by each dot is the quantity of the bid (in £millions). The three bids made by the bidder described above are highlighted in bold.

---

[13] A bidder can, of course, restrict each of its bids to a single variety. Note also that a bidder who wants to guarantee winning a fixed total quantity can do so by making a bid at an arbitrarily large price for its preferred variety, and at an appropriate discount from this price for the other variety. We discuss other kinds of bids that bidders can be permitted to make in the next section.

[14] Since all bidders prefer weak to strong, the plots all fall below the 45º line -- this special feature of our example is of no importance to the auction design.

If, for example, the Bank wishes to loan £2.5 billion, and there are a total of £5.5 billion in bids, then the possible sets of *excluded* bids are any set of bids included in a rectangle drawn with two sides being along the axes and encompassing £3 billion in bids. Figure 1a shows one possible rectangle of rejected bids, bounded by the vertical line at 5.92% (for weak) and the horizontal line at 5.65% (for strong). Bids outside the rectangle are accepted.

Those bids for which both offers exceed the corresponding cut-off prices (that is, those bids to the north-east of the rectangle) are allocated to the variety for which the cut-off price is further below the offer. Thus bids that are both north of the rectangle, and north-west of the diagonal 45° line drawn up from the upper-right corner of the rectangle, receive strong; the other accepted bids receive weak.

The Bank uses a uniform-pricing rule for each variety. So all bids accepted for strong pay the same minimum (cut-off) price for strong (5.65% in the example illustrated), and all bids accepted for weak pay the same minimum price for weak (5.92% here).[15]

Of course, Figure 1a shows only one of many points that represents a pair of cut-off prices that rejects exactly £3 billion of bids. There is generally exactly one such point on any 45° line on the plane[16] -- the set of all these points is the stepped downward–sloping line in Figure 1b.

Every possible price pair on Figure 1b's stepped line implies both a price-difference and (by summing the accepted bids below the corresponding 45° line) a proportion of sales that are weak. As the price-difference is increased, the proportion of weak sales decreases. Using this information we can construct the downward-sloping "demand curve" in Figure 2. (Note that the axes of Figure 2 are different from those of Figure 1!)

---

[15] We discuss the possibility of discriminatory pricing in Appendix 2 of Klemperer (2008).

[16] Moving north-east along any 45° line represents increasing all prices at the same rate (that is, with a constant price difference between varieties). Because the marginal bid(s) is usually rationed, there is usually a single critical point that rejects the correct volume of bids. But if exactly £3 billion of bids can be rejected by rejecting entire bids, there is generally an interval of points between the last rejected and the first accepted bid. (We would then choose the most south-westerly of these points, though other selection rules are possible.)

Figure 1b. Feasible Pairs of Cutoff Rates

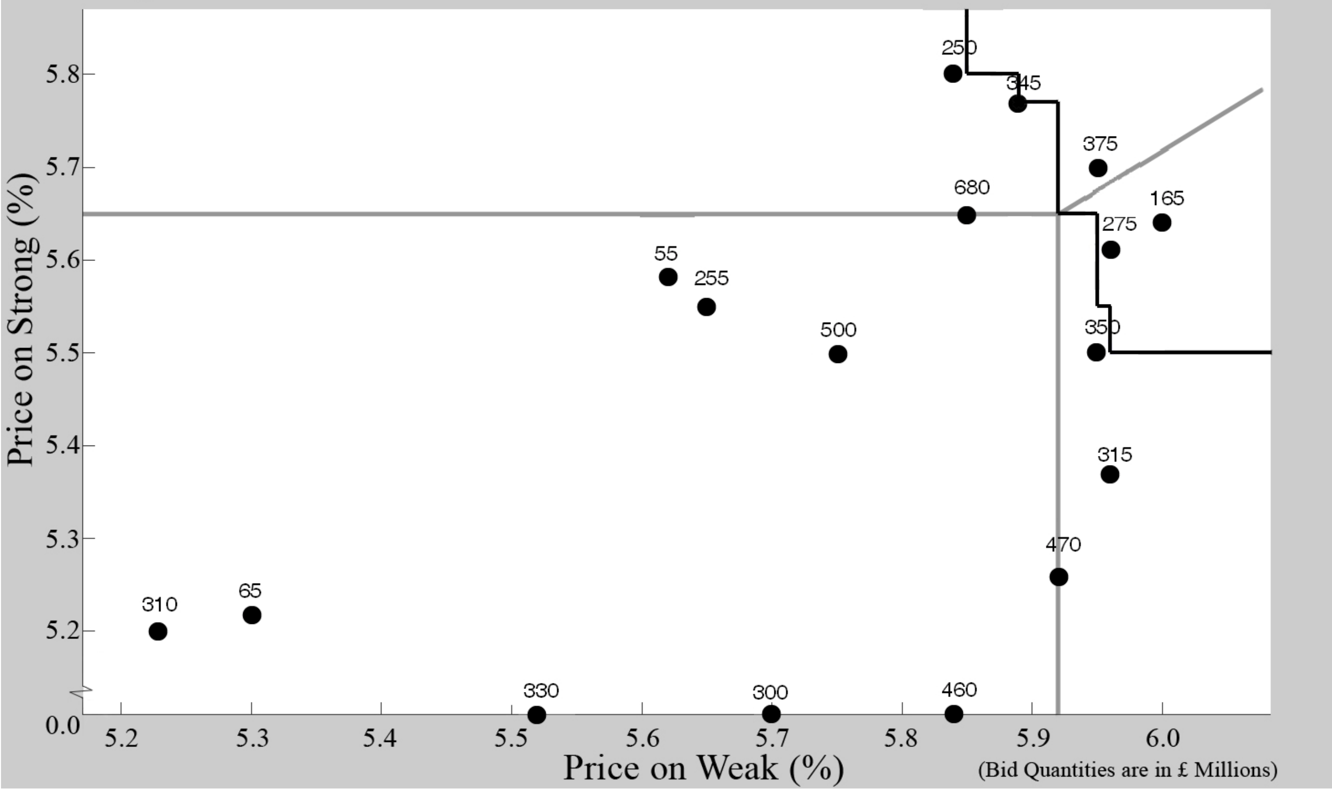

Figure 2. Demand and Supply for Weak Collateral

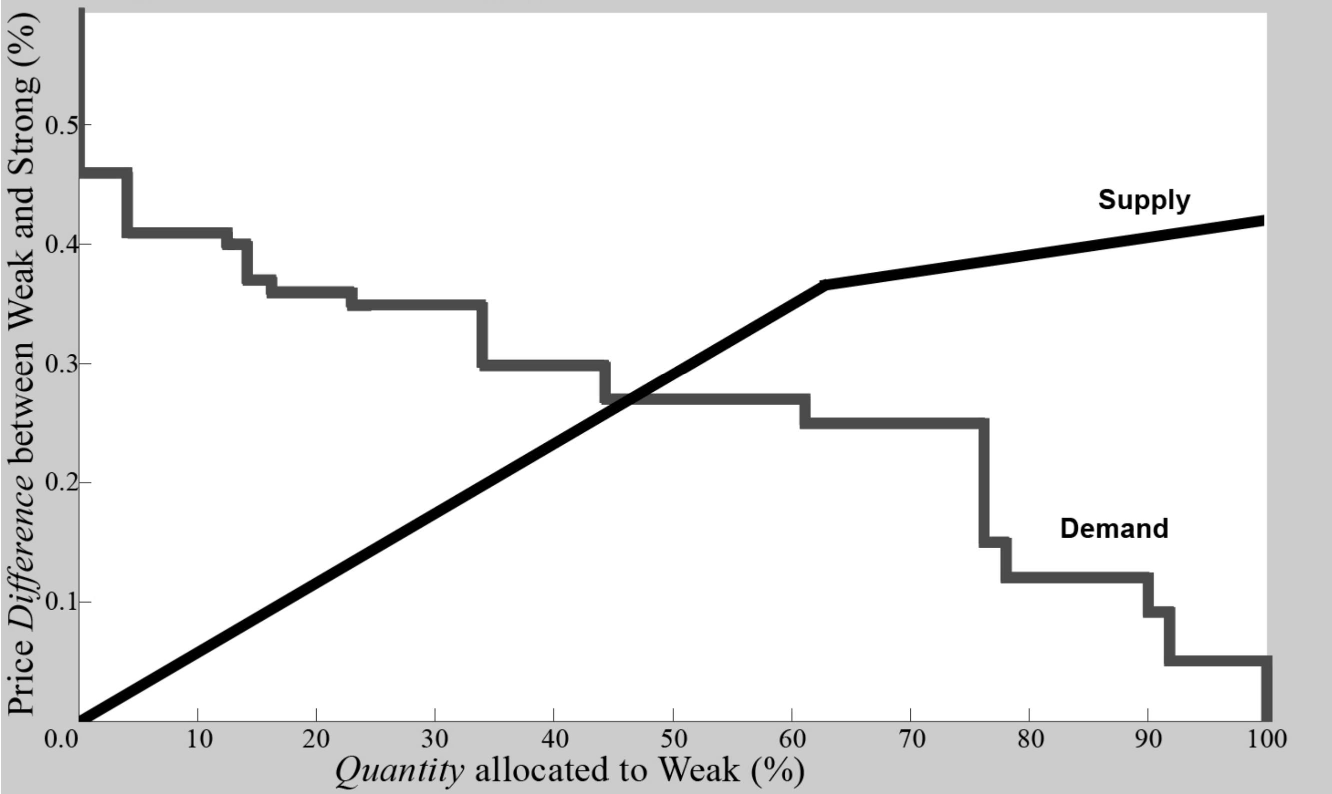

The Bank could give itself discretion to choose any point on the "demand curve" (equivalently, any feasible rectangle in Figures 1a, 1b) after seeing the bids.

Alternatively, the Bank could precommit to a rule that will determine its choice, that is, precommit to a "supply curve" or "supply schedule" such as the upward-sloping line in Figure 2.[17] Choosing a horizontal line would fix the price differential between weak and strong. (Equivalently, the rectangle selected from the feasible set whose upper-right corners are shown in Figure 1b would be constrained to be a fixed amount wider than it was high.) Choosing a vertical line would fix the proportions of the two goods. Choosing an upward sloping schedule, as illustrated, so the proportion allocated to weak increases with its price-premium, corresponds to "an average of" these approaches.[18]

The point of intersection between the Bank's supply curve and the "demand curve" constructed from the bids determines the price differential and the percentage of weak sold in the auction. With the supply schedule illustrated in Figure 2, the price difference is 0.27% and the proportion of weak is 45% -- corresponding to accepting the bids outside the rectangle shown in Figure 1a, at the prices at the rectangle's "northeast" corner (5.92% for weak, and 5.65% for strong).[19]

The design generalises easily: [20]

---

[17] If the auctioneer predetermines its supply curve, it can choose whether or not to preannounce it. Maintaining the flexibility to alter it mitigates bidders' market power (see Klemperer and Meyer (1989), McAdams (2007), etc.).

[18] One proposal for the U.S. TARP was that the government should spend a predetermined amount on each class of security; this corresponds to choosing the multi-dimensional equivalent of a vertical supply schedule. Another proposal was to develop a "reference price" for each asset, and buy those that were offered at prices furthest below the reference prices; this corresponds to choosing a horizontal supply schedule. As we noted above, both these suggestions were flawed. Choosing an upward-sloping supply schedule maintains the advantage of the reference-price approach, while limiting the costs of mispricing, (Choosing the optimal supply-schedule slope involves issues akin to those discussed in Weitzman (1974), Klemperer and Meyer (1986), Poole (1970), etc.)

[19] Note that by specifying the proportion of weak, the intersection of the curves in Figure 2 generally determines what fractions of any bids on the borders of the rectangle are filled, and how any bids on the 45° line are allocated between the goods.

[20] For a problem like the Bank of England's, the features described thus far may suffice. The U.S. TARP might have benefited from some of the next section's extensions, including quantity constraints on groups of bids, and combination bids that guarantee winning minimum amounts, but win more if prices are favourable.

3. Easy Extensions

3.1 Multiple buyers and multiple sellers

It is easy to include additional potential sellers (i.e., additional lenders of funds, in our example). Simply add any additional seller(s)' maximum supply to the total that the auctioneer sells, but allow them to participate in the auction as usual. If a potential seller wins nothing in the auction, the auctioneer has sold the seller's supply for it. If the potential seller wins its total supply back, there is no change in its position.

3.2 "Swappers" who might want to be on either side of the market

Exactly the same approach permits a trader to be on either, or both, sides of the market. If, for example, letting the auctioneer offer its current holdings of weak, a bidder in the auction wins the same amount of strong, it has simply swapped goods (paying the difference in the market-clearing price).

3.3 Other easy extensions

Several other extensions are also easy. For example, bidders can be allowed to ask for different amounts of the different goods in a bid. Or a bidder can specify that a total quantity constraint applies across a group of bids. And there can, of course, be more than two goods, with a cut-off price for each, and a bid rejected only if *all* its offers were below the corresponding cut-off prices.

Bidders can express more complex preferences by using several bids in combination: for example, a bidder might be interested in £100 million weak at up to 7%, and £80 million strong at up to 5%. However, even if the prices of funds are high, the bidder wants an absolute minimum of £40 million. This can be implemented by making all of the following four bids, if negative bids are permitted:

1. £40 million of {weak at maximum permitted bid OR strong at maximum permitted bid *less* 2%}.[21]
2. £100 million of weak at 7%.[22]
3. £80 million of strong at 5%.
4. *minus* £40 million of {weak at 7% OR strong at 5%}.

The point is that the fourth bid kicks in at exactly the same point where one of the second and third bids are accepted, and this negative bid then exactly cancels the first bid for £40 million "at any price" (since 2% = 7% - 5%).

Other extensions allow bidders to express still richer forms of preferences, but there are limits, as we now discuss.

4. Further Extensions, and the Relationship to the SMRA

To see further potential extensions, observe that our auction is equivalent to a static (sealed-bid) implementation of a simplified version of a two-sided simultaneous multiple round auction (SMRA).

Begin by considering the case in which the auctioneer has predetermined the quantity of each variety it wishes to offer, for example, £7 billion of strong and £3 billion of weak, and the bids represent bidders' true preferences.[23] Then in a standard SMRA, if each bidder bids at every step of the SMRA to maximise its profits at the current prices given these preferences, the outcome will be exactly the same as that of our procedure (in the limit as the bid increments are zero).[24] In other words, when bidders behave

[21] The maximum permitted bid in the auction is set at a high enough level that it will be used only by bidders who want to guarantee winning at any price.
[22] Formally, this bid includes "OR of strong at 0%", and the next bid includes "OR of weak at 0%".
[23] For example, if a bidder's most north-easterly bid is "£100 million of {weak at 7% OR strong at 5%}" then the bidder is indifferent between receiving any of (i) nothing (ii) £100 million of weak at 7%, and (iii) £100 million of strong at 5%, etc. (see Klemperer (2008)).
[24] That is, imagine the bidders take turns to make bids in many ascending auctions that are run simultaneously (in this example, 7 billion auctions for a single £1 of strong, and 3 billion auctions for a single £1 of weak). When it is a bidder's turn, it can make any new bids it wishes provided they beat any existing winning bid of a competitor by some arbitrarily small bidding increment. (The bidder cannot top up or withdraw any of its own existing bids, but it would anyway not want to do so.) The bidder makes the bids that would maximize its profit if all the auctions were to close immediately after its bids. The bidders continue to take turns until no one wants to submit any new bids. Then the outcome will be the same as ours.

competitively as price takers, as is rational behaviour when the number of bidders is not too small,[25] both our mechanism and the SMRA (in the limit as the bid increments are zero) simply select the competitive-equilibrium price vector. (When this price vector is not unique, they both select the unique vector among these that is lowest in every element).[26]

The case in which the auctioneer offers a supply curve relating the proportions of the different varieties sold to the price differences, rather than simply fixing the quantity of each variety, is not much harder: we now think of the auctioneer as acting *both* as the bid-taker selling the maximum possible quantity of both varieties, *and* as an additional buyer bidding to buy units back to achieve a point on its supply curve. That is, in our example in which the Bank auctions £10 billion, we consider an SMRA which supplies £10 billion weak *and* £10 billion strong, and we think of the Bank as an additional bidder who has an inelastic total demand for £10 billion and who bids in exactly the same way as any other bidder. (That is, whenever it is the Bank's turn to bid, it will bid on one or both varieties to both restore its quantity of winning bids to £10 billion and win the quantity of each variety that puts it back on its supply curve, given the interest rate-differential it faces.[27])

More generally, if there are other sellers (or "swappers") we consider an SMRA in which their potential sales (or "swaps") are added to those offered in the auction, and think of these participants as bidding for positive amounts like any other bidders.

---

[25] We assume bidders would not update their preferences in response to observing others' bidding. That is, they have "private values".

[26] To understand this, recall that Milgrom (2000, theorem 3) shows that when goods are mutual substitutes in demand for all bidders, the SMRA with infinitesimal bidding increments selects the (minimum price) competitive equilibrium. (See also Crawford and Knoer (1981, theorem 3) and Kelso and Crawford (1982, theorem 4).) Furthermore, any outcome of our procedure is a competitive equilibrium (since each bidder takes its preferred allocation given the prices), and the preferences represented by the types of bids we have discussed imply substitutes demand.

[27] The prices on winning bids on any variety will differ by most one bid increment. Ignoring this difference, the Bank can always return to its supply curve, since the weak-minus-strong price difference can only be more (less) than when it previously bid if its weak (strong) bids have all been topped, so it can increase the quantity of strong (weak) it repurchases relative to its previous bids, as it will wish to do in this case. (We assume its supply curve also exhibits substitutes preferences. In this example, that simply means that the quantity of weak it supplies (equivalently, the quantity of strong it repurchases) is increasing in the weak-strong price difference.)

So our procedure is equivalent to one in which bidders submit their preferences, and the auctioneer and other (potential) sellers submit their supply curves, and a computer then calculates the equilibrium of a SMRA. (Note that though the way we described the auctioneer's supply function may have obscured this, our procedure is symmetric between buyers and sellers.[28])

The only difference between our procedure and a (two-sided) SMRA with "proxies" (bidding rules), is that we have limited the preferences that bidders' proxies can express. In principle we could allow bidders to specify *any* proxy bidding rule, subject to computational issues. These issues are not very challenging in our simple example (or in the Bank of England's actual problem).

However, some bidding rules can lead to different outcomes depending upon the order in which bidders take turns to bid. Consider, for example, a bidding rule that attempted to maximise a bidder's surplus subject to the constraints that (total quantity purchased of varieties A and B $\leq \alpha$) and (total quantity purchased of varieties B and C $\leq \beta$). Then an increase in A's price might make a bidder want to bid for more B instead of A, and this could make it want to bid for less C, so it might now regret an earlier bid for C -- and indeed it might never have made the earlier bid if the order of play had resulted in prices evolving differently.[29] So some constraints on bidding rules are probably desirable.[30]

In fact, limiting bidding rules to correspond to preferences that cannot be revised in response to others' bids makes exercising market power much harder than in a standard SMRA.[31] If the number of bidders is small, they will bid less than their true demands in

---

[28] The implications of an auctioneer (such as the Bank of England) being a single player and having non-profit objectives are discussed in Appendix 1 of Klemperer (2008). Our procedure is not quite symmetric if the auctioneer doesn't precommit to its supply schedule, but if bidders behave competitively their bids are unaffected by this schedule's position, and hence also by when it is chosen.

[29] The problem is the standard one that competitive equilibrium might not exist without "substitutes preferences" (see Milgrom 2000, theorems 3, 4; Kelso and Crawford (1982, pp. 1502-4)). In independent work, Milgrom (forthcoming) explores how to restrict bidders' expressions of preferences to avoid such problems.

[30] Thinking about what preferences a bidder may wish to express in a SMRA suggests possible extensions to the permitted forms of bidding discussed above (see Klemperer 2008, Appendix 3).

[31] Limiting bidders' strategies in this way may be less desirable if "common-value" components to valuations mean bidders might justifiably revise their references after observing others' bidding.

any procedure that charges them constant per-unit prices,[32] but it is much harder for them to (ab)use their market power in our mechanism because they have to submit their bids before seeing any information about other bidders' demands. And implicit collusion (coordinated demand reduction) and predatory strategies are hard or impossible with a single simultaneous bidding round. In a standard dynamic SMRA, by contrast, bidders can learn from the bidding when and where market-power-abusing strategies are likely to be profitable, and they can make bids that signal threats and offers to other bidders and easily punish those who fail to cooperate with them.[33]

Finally, the parallel with standard sealed-bid auctions makes our mechanism more familiar and natural than the SMRA to counterparties. In contexts like that of the Bank of England, our procedure is much simpler to understand than an SMRA.

The advantages of our mechanism over running multiple separate auctions were outlined in the Introduction; Klemperer (2008) gives many more details of the comparison of our mechanism with alternatives.

5. Conclusion

I have described a simple-to-use, sealed-bid, auction that allows bidders to bid on multiple assets simultaneously, and bid-takers to choose supply functions across assets. It can be used in environments in which a simultaneous multiple-round auction (SMRA) is infeasible because of transaction costs, or the time required to run it. The design also seems more familiar and natural than the SMRA to bidders in many potential applications, and makes it harder for bidders to exercise market power. Relative to

---

[32] A proxy bidder that always bids to maximise its profits at the current prices does not take into account the impact of its demand for an incremental unit on the price it will pay for its inframarginal units, but bidders can still exercise market power by submitting multiple proxies, i.e. multiple bids, with all but their first bids at prices below their true values.

[33] In a standard SMRA, a bidder can follow "collusive" strategies such as "I will bid for (only) half the lots if my competitor does also, but I will bid for more lots if my competitor does". See, e.g., Klemperer (2002, 2004). In our procedure, bidders can specify preferences for objects, but these preferences cannot be functions of others' bids.

Note also that even without deliberate predation, an SMRA may discourage entry of bidders who feel less able than their rivals to use the information learned between rounds.

running separate auctions for separate goods, it yields better "matching" between suppliers and demanders, reduced market power, greater volume and liquidity, and therefore also improved efficiency, revenue, and quality of information. Its potential applications therefore extend well beyond the financial contexts for which I developed it.[34]

References

Ausubel, Lawrence, and Peter Cramton (2008). "A Troubled Asset Reverse Auction." Working Paper, University of Maryland.

Ausubel, Lawrence, Peter Cramton, Emel Filiz-Ozbay, Nathaniel Higgins, Erkut Ozbay, and Andrew Stocking (2008). "Common-Value Auctions with Liquidity Needs: An Experimental Test of a Troubled Assets Reverse Auction." Working Paper, University of Maryland, December 2008.

Binmore, Ken and Paul Klemperer (2002). "The Biggest Auction Ever: the Sale of the British 3G Telecom Licenses", *Economic Journal*, 112, C74-C96.

Crawford, Vincent P., and Elsie Marie Knoer (1981). "Job Matching with Heterogeneous Firms and Workers." *Econometrica*, 49, 437-450, March 1981.

Kelso, Alexander S. Jr., and Vincent P. Crawford (1982). "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50, 1483-1504, November.

Klemperer, Paul (2002). "What Really Matters in Auction Design." *Journal of Economic Perspectives*, 16 (1), 169-189.

Klemperer, Paul (2004). *Auctions: Theory and Practice*, Princeton University Press, Princeton, US.

Klemperer, Paul (2008). "A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions". Working paper, Oxford University.

Klemperer, Paul and Margaret Meyer (1986). "Price Competition vs. Quantity Competition: The Role of Uncertainty." *Rand Journal of Economics*, 17 (4), 618-638.

Klemperer, Paul and Margaret Meyer (1989). "Supply Function Equilibria in Oligopoly under Uncertainty," *Econometrica*, 57, 1243-1277.

McAdams, David (2007). "Uniform-Price Auctions with Adjustable Supply." *Economics Letters*, 95, 48-53.

Milgrom, Paul (2000). "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 108, 245-272.

Milgrom, Paul (forthcoming). "Assignment Messages and Exchanges." Working paper, Stanford University.

Poole, William (1970). "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model." *The Quarterly Journal of Economics*, 84 (2), 197-216, May.

Weitzman, Martin (1974). "Prices vs. Quantities," *Review of Economic Studies*, October.

[34] At the time of writing a Latin American regulator is considering a proposal (not by me) to use it to purchase two close-substitute "types" of electricity.