This is the version circulated in 2008; a revision was published in the Journal of the European Economic Association in 2010; a larger revision is at www.nuff.ox.ac.uk/users/klemperer/productmix.pdf

A New Auction for Substitutes:

Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions

Paul Klemperer, 2008

INCOMPLETE (NON-CONFIDENTIAL) VERSION

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

Abstract

I describe an auction designed in late 2007- early 2008 to help the Bank of England fight the credit crunch.

A similar design would have been useful, and might have been used, if the U.S. Treasury had pursued its original plan to spend much of its $700 billion TARP funding to buy distressed mortgage-backed securities.

The same design is effective for auctioning multiple substitute goods whenever multiple round auctions are infeasible.

It is a simple-to-use, static (sealed-bid) mechanism. But like a two-sided simultaneous multiple round auction, it permits bidders to bid on multiple assets simultaneously, and bid-takers to choose supply functions across assets. So bids for different assets are forced into competition with each other.

The design therefore yields greater volume, greater efficiency, better information, and more revenue, than running multiple static (sealed-bid) auctions.

Disclaimer: **I have been an adviser to the Bank of England since late 2007, and I advised the U.S. Treasury from September to November 2008. I have also recently consulted with other Central Banks, government agencies, and other organisations about related issues.**

The views here are personal and do not represent the views of these, or any other, organisations.

Although I use a generic Central Bank's problem as an example, none of the discussion refers specifically to the Bank of England or any other actual Central Bank, except where specifically noted.

Some sensitive information, and Appendices about detailed design issues, have been removed from this version.

Acknowledgements: I am especially grateful to Jeremy Bulow …. also to officials of several Central Banks and to .... [to be completed later]

Nuffield College, Oxford, OX1 1NF, England; paul.klemperer@economics.ox.ac.uk
Tel: +44 777 623 0123; fax: +44 1865 511749 © Paul Klemperer 2008

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. The Bank urgently wanted to supply liquidity to banks, and was willing to accept a wider range of collateral than it had traditionally accepted if that was necessary to lend the desired amount. But it wanted a correspondingly higher interest rate against weaker collateral. Furthermore, because financial markets move fast, any auction had to take place at a single instant -- a multi-stage auction was ruled out because bidders who had entered the highest bids early on might change their minds about wanting to be winners before the auction closed,[1] and because 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'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.[2] (The multiple 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 subprime mortgage-backed securities. As above, the volatility of financial markets, and their sensitivity to news, would have made a multi-round auction problematic.[3] And because there were a large number of closely-related but differentiated assets, some with very concentrated

[1] Confirmatory evidence is that most bids in the actual sealed-bid auctions that were run were made in the last few seconds. For a multi-round auction to have any merit, untopped bids cannot be withdrawn without incurring penalties.

[2] Put differently, 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.

[3] The potential feedback effects between the financial markets and any dynamic auction seem especially severe in this context.

ownership, an auction in which the buyer simply pre-specified the quantity of each type of security to purchase would not have ensured adequate competition.

This note outlines a solution to all these problems.[4] In early 2008, I proposed a version of it to the Bank of England (which is holding consultations on the proposal),[5] and in autumn 2008, I and others made a similar proposal to the U.S. Treasury (which might have adopted a related design if it had not abandoned its plans to buy subprime assets).[6]

We begin in section 1 by describing the problem of a Central Bank (such as the Bank of England[7]) in very general terms. Section 2 outlines my solution, and section 3 provides a simple graphical illustration for the two-good case. The generalization to many goods is straightforward. Section 4 develops some other easy extensions that are probably unnecessary in Central Banking, but may be valuable elsewhere; they might, for example, have been valuable in the U.S. TARP.[8] We observe in section 5 that our design is essentially a two-sided "proxy" implementation of a standard simultaneous multiple

[4] Subsequent to proposing this solution to the Bank of England (see next note) I learned that Paul Milgrom was independently pursuing related ideas. He and I therefore worked together for the U.S. Treasury (see next note but one). Milgrom (2008) shows how to represent a wide range of bidders' preferences very elegantly while at the same time restricting to substitutable preferences, and goes further than the current paper by showing that his highly efficient linear-programming approach yields an integer allocation when demands and constraints are integer -- this property could be very important in some applications, even though it is not in a context such as that of the Bank of England, for which my proposal seems more straightforward and transparent.

[5] The Bank consulted me a few months after the credit crunch began. The crisis began in early August 2007, and a bank run resulted in Northern Rock's collapse in mid-September. In late September and the first half of October, the Bank of England ran four auctions to supply additional liquidity to banks -- but no bids were received in any of them (for reasons that are not the subject of this note). The Bank consulted me shortly afterwards, and I got assistance from Jeremy Bulow and Daniel Marszalec. Starting in December, the Bank ran additional simple (more successful) auctions (see below) while developing and considering the ideas discussed here. There was essentially no change in our ideas after February 2008, but the continued unsettled financial markets, and the fact that the simple auctions that began in December 2007 were achieving the Bank's main objectives, meant the consultation process with counterparties etc. began only in October 2008.

[6] Jeremy Bulow, Jon Levin, Paul Milgrom, and I made a joint proposal to the U.S. Treasury. Other consultants, too, proposed a static (sealed-bid) design, and it is likely that a sealed-bid design would have been used, although some advisers, including Ausubel and Cramton (2008), argued that a simultaneous multi-round auction was viable in spite of the difficulties described above.

[7] *Note we do NOT give full details of the Bank of England's specific objectives and constraints here. Not all of the Bank of England's concerns are discussed, and some of the issues we do discuss are of little or no significance to the Bank of England. Furthermore, the general solution I describe contains far more features than are likely to be required in a Central Banking application.*

[8] *(Only) limited discussion of the TARP is included below; it is hoped to include more detailed discussion in a future draft.*

round auction (SMRA),[9] and then discuss further extensions. Section 6 contrasts the merits of our approach with alternatives, and section 7 concludes.

1. The Problem

We consider a Central Bank (henceforth "the Bank") that wants to lend a given amount, and prefers to do so against higher-quality collateral and for higher interest rates. The auction has to take place at a single moment in time.

The most straightforward approach (and the one adopted by the Bank of England pending development of our proposals) is to run separate sealed-bid auctions for high-quality and for low-quality collateral.[10, 11]

Of course, this approach has the crucial problem that the Bank is forced to choose how much to offer against each collateral before learning bidders' preferences. Furthermore, bidders would like to find out the difference between the different auctions' clearing-prices before bidding, but cannot do this and have instead to make guesses about which auction will offer them best value.[12] Outcomes are therefore erratic and inefficient: funds

[9] See Milgrom (2000) for a description of the SMRA. SMRAs are commonly used for auctioning radio spectrum see, for example, Binmore and Klemperer (2000). For general introductions to the economics literature on auctions, see Klemperer (1999, 2002, 2004), Krishna (2002), Milgrom (2004), and Menezes and Monteiro (2005).

[10] See note 5, above. In the case of the Bank of England, what I call the "low-quality collateral" auctions actually permitted both high and low quality collateral, but did not discriminate between bids on the basis of the collateral offered, either in selecting the winning bids, or in the interest rates paid. So, of course, bidders had strong incentives to offer only the lowest quality collateral permitted. Although the detailed results are confidential, it can be inferred from tables in Bank of England publications that at least a large part (and perhaps essentially all) of the collateral offered in any auction was of the lowest permitted quality.

[11] Most of the existing debate had been about whether such auctions should use discriminatory (pay-as-bid) or uniform pricing. In this context of separate sealed-bid auctions, a reasonable case could be made for either. Indeed, the Bank of England used discriminatory auctions, while the US Federal Reserve used uniform-price auctions for its Term Auction Facility (which performed a similar role in the US to the Bank of England's auctions, see Armantier et al, 2008), and the Bank of Canada used both kinds of pricing in its similar auctions.

[12] Consider, for example, a counterparty who wants, say, £300 million. Should it bid for £300 million against each type of collateral and risk being allotted £600 million? Or to avoid that risk, should it bid for £300 million against one type of collateral but not the other? It might then see money allocated against the other type of collateral, at a rate it would have been willing to pay. Or should it bid for £150 million in each? Whatever it does, it may after the fact regret borrowing on the wrong type of collateral given the

are unlikely to go to those who value them most, and those bidders who do win them might be inefficiently allocated across collaterals.

Furthermore, when funds against separate collaterals are auctioned separately, offers made against one collateral provide no competitive discipline on the offers made against other collaterals. So each individual auction is more sensitive to market power, to manipulation, and to informational asymmetries, than if all bidders' offers competed directly with each other in a single auction. Interest rates (i.e., bid-taker revenues) are correspondingly generally lower.

These problems also reduce the auctions' value as a source of information for the Bank and other market participants. The same problems may also reduce participation in the auctions, which creates "second-round" feedback effects which further magnify the problems.

In short, the straightforward approach generates poor outcomes for both the bidders and the bid-taker.

Our solution addresses these problems.

2. The Solution

Our proposal is straightforward in concept: allow each counterparty (bidder) to offer one or more sets of bids; each set contains an interest-rate bid for one or more collaterals, and the bids in each set are mutually exclusive. The bid-taker (the Bank) looks at all the sets of bids and then chooses its preferred interest rates (a separate, uniform, rate for each collateral), according to some predetermined (but not necessarily preannounced[13]) rule.

---

market clearing rates, and/or borrowing too much or too little given the market clearing rates. (The problem is independent of whether uniform or discriminatory pricing is used.)

[13] A Central Bank has enough institutional credibility that it would not be expected to behave strategically if it did not pre-announce its rule.

From each set of bids offered by each bidder, the Bank accepts the one that gives the bidder the greatest surplus evaluated at those interest rates[14] (or no bid if all bids yielded negative surplus).

The idea is to allow the Bank to look at demand *before* choosing how much to offer against each collateral, at the same time as also allowing each bidder to achieve its best possible outcome given the interest rates the Bank actually selects. (By making contingent bids, bidders can in effect decide how much, and against which collateral, to borrow *after* seeing the interest rates chosen.)

The question, of course, is whether this can actually be implemented, and whether it can be done in a way that is simple and robust, and also sufficiently easy for bidders to understand that they are happy to participate. We now show that this is feasible. We begin by illustrating a simple approach similar to the one I proposed to the Bank of England, before discussing the range of possibilities.

3. A Simple Two-Good Example

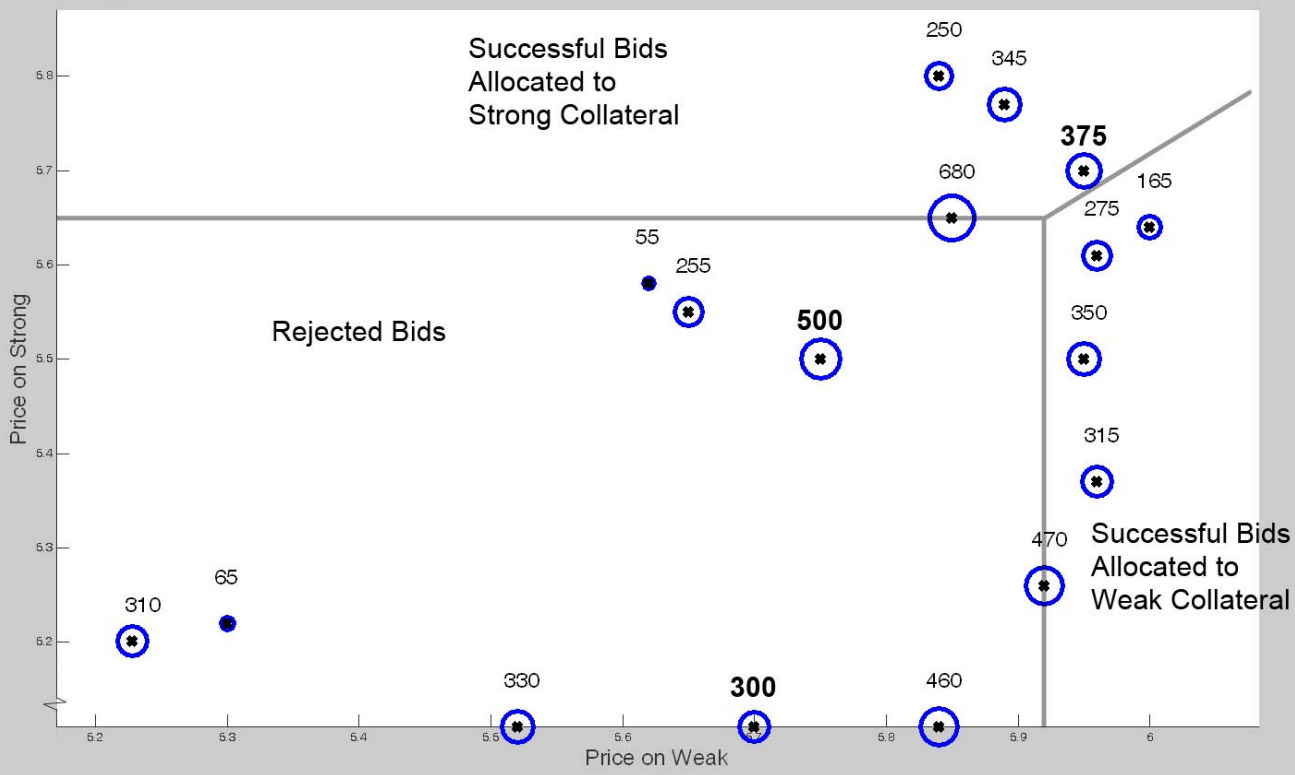
Suppose there are just two classes of collateral, "strong" and "weak".[15] Each bidder (counterparty) is allowed to make a number of bids in the auction. Each bid is for an amount of money, and includes two interest rates. One rate is that which the counterparty is willing to pay if it borrows against strong collateral; and the other relates to borrowing against weak collateral. The two outcomes would be mutually exclusive. If a counterparty has, or wishes to use, only one type of collateral, it is allowed to bid a zero interest rate for the other type of collateral, which it cannot or will not use -- that will guarantee that the undesired collateral will never be selected from that bid.

[14] If bids are tied, the Bank can choose which bid to accept. If the highest-surplus bid(s) yield zero surplus, the Bank can choose what fraction(s) to accept.

[15] "Strong" might correspond to the “OMO" or "ordinary" collateral the Bank of England traditionally accepted in its "open market operations". "Weak" might correspond to the "wider" or "extended" collateral that the Bank of England was willing to lend against in the stressed circumstances that developed from autumn 2007.

Thus, for example, a bidder might make a bid to borrow £375 million at 5.95% if it were allowed to use weak collateral and 5.7% if it were required to use strong collateral. It might make a second bid to borrow an additional £500 million at 5.75% if it could use weak collateral and 5.5% if it had to use strong collateral. It might make a third bid to borrow £300 million against weak collateral at 5.7% and 0% against strong collateral (which would be read as a refusal to borrow against strong collateral as part of this particular bid). Each of these bids would be interpreted as either/or bids in the sense that the Bank would accept at most one of the two offers, namely the one that gives the borrower the better deal in terms of the difference between the rate the borrower offers and the market-clearing rate -- this prevents the bidder from having to worry that one of its bids will defeat another that it would have preferred to have had accepted.[16]

Figure 1a. A Possible Allocation of Funds



An example of the universe of bids submitted by all the bidders is illustrated in Figure 1a above. Strong-collateral rates are plotted vertically, and weak-collateral rates

[16] We discuss other kinds of bids that bidders can make in the next section.

horizontally, so that each dot in the chart represents a paired bid.[17] (The number by each dot is the amount of the bid in millions of pounds.) The three bids described in the previous paragraph are numbered in bold. (Observe that the vertical axis of the figure is "broken" so that the bids for £330m, £300m, and £460m are all for 0% on strong -- i.e., these bids are equivalent to traditional "unpaired" bids on weak collateral only.)

If, for example, the Bank wishes to lend £2.5 billion, and there are a total of £5.5 billion in either/or bids, then £3 billion of bids have to be refused. Exactly which £3 billion would be refused would be determined by the rule that the Bank chose to use in providing funds. The possible sets of excluded bids would be any set of bids included in a rectangle drawn with two sides being along the axes and encompassing £3 billion in bids. Each possible rectangle is uniquely identified by the pair of interest rates corresponding to the upper right corner of the rectangle; bids above ***either*** of these "cut-off" interest rates are accepted, while bids below ***both*** cut-offs are rejected. Figure 1a shows one possible pair of cut-off rates, given by the vertical line at 5.92% (for weak collateral) and the horizontal line at 5.65% (for strong collateral). Bids within the rectangle are rejected, and the others are accepted.

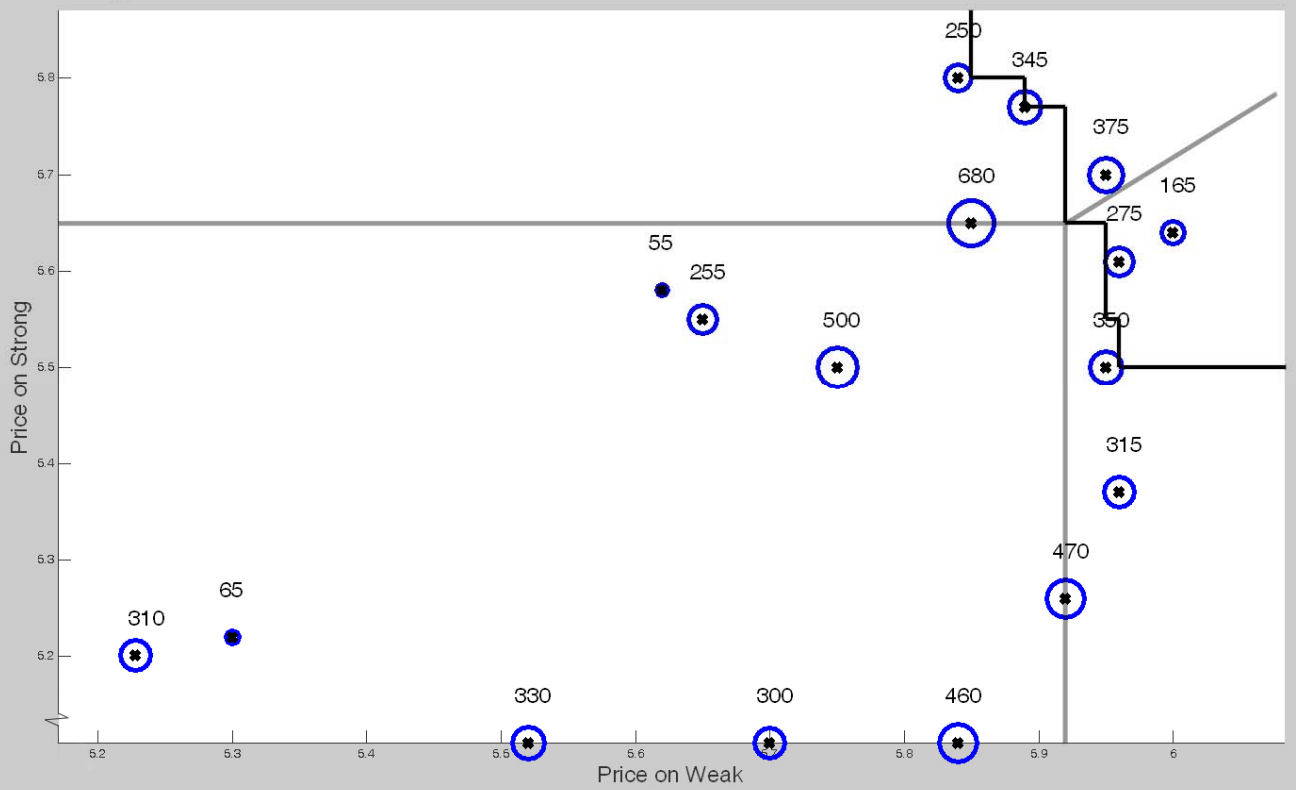
Those bids for which both offers exceed the corresponding cut-off rates (that is, those bids to the north-east of the rectangle) are allocated to the collateral for which the cut-off rate is further below the offer. Thus bids that are both north of the rectangle, and north-west of the diagonal 45° line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

The Bank uses a uniform-pricing rule for each collateral. So all bids accepted against strong collateral pay the same minimum (cut-off) interest-rate for strong collateral, and all bids accepted against weak collateral pay the weak-collateral interest-rate. (We discuss the possibility of discriminatory pricing in Appendix 2.)

[17] Since weak-bids are all greater than strong-bids, the plots all fall below the 45º line -- this is a special feature of our Central Banking example, and is of no importance to the auction design.

Of course, Figure 1a shows only one of a continuum of possible pairs of cut-off rates that would reject exactly £3 billion of bids. If we draw a 45° line through any point on the graph at which the difference between interest rates is not too extreme, there will generally be a single point on that 45° line at which exactly £3 billion of bids are rejected.[18] (If there is more than one point we choose the most south-westerly, though other selection rules are possible.) As we move the 45° line southeast (relatively higher interest rates on weak collateral) the critical point representing the pair of cut-off rates moves either down or to the right. All of the possible pairs are joined by the stepped downward–sloping line in Figure 1b.

Figure 1b. Feasible Pairs of Cutoff Rates



Every possible cut-off pair on Figure 1b's stepped line implies both a difference in interest rates and (by summing the accepted bids below the corresponding 45° line) a

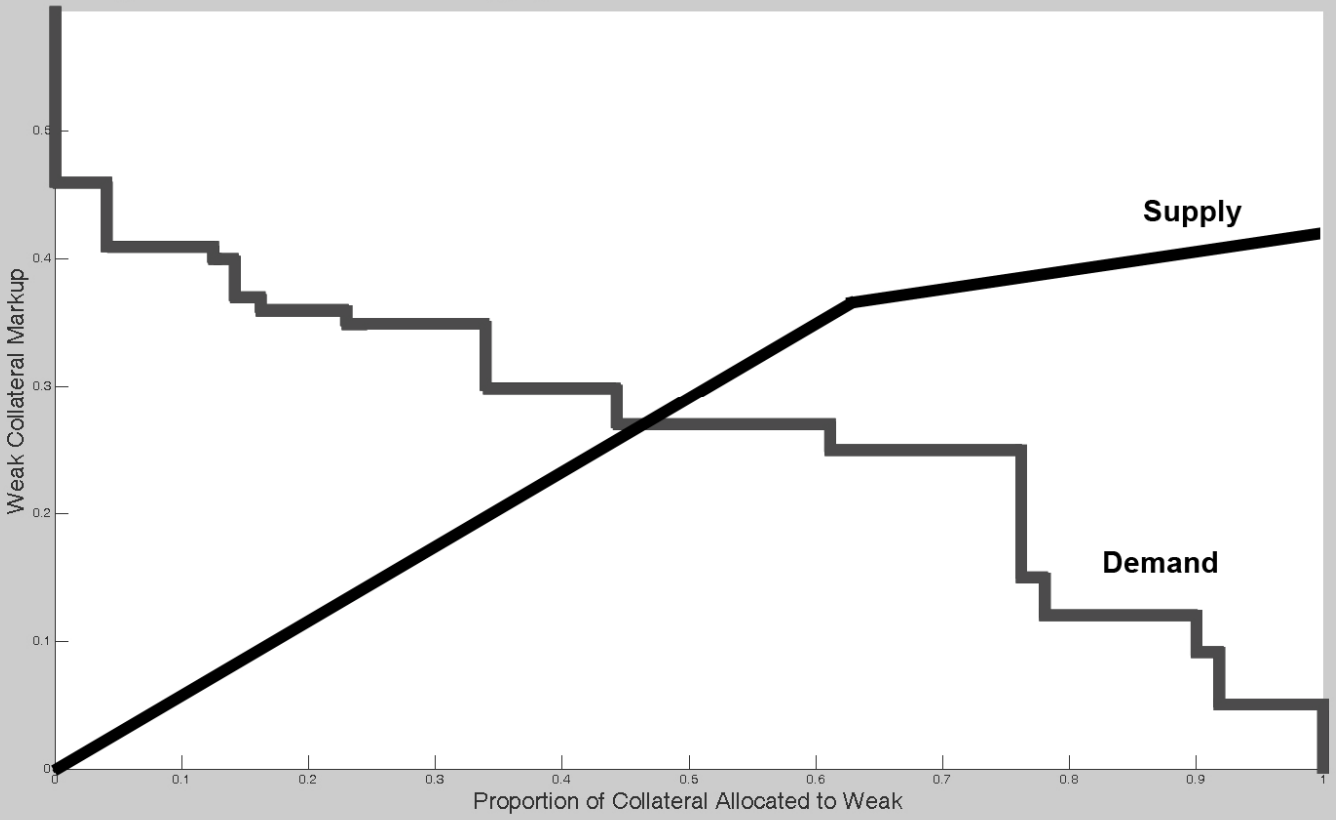
[18] If exactly £3 billion of bids can be rejected by rejecting entire bids (equivalently, the amount to be accepted can be made up with entire bids), there will generally be an interval between the last rejected bid and the first accepted bid. Usually, however, the marginal bid(s) will be rationed, so supply will equal demand at only one point on any 45° line.

proportion of funds allocated to weak collateral. As the interest-rate difference is increased, the proportion allocated to weak collateral decreases. Using this information we can construct the downward-sloping “demand curve” (the stepped line) in Figure 2.

(Note that the axes of Figure 2 are different from those of Figure 1.)

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.

Figure 2. Demand and Supply for Weak Collateral



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. If the Bank chooses the supply curve illustrated, it would allocate no funds to weak collateral if all the interest rates offered were the same, but the proportion it would allocate to weak collateral would increase with the differential between the two interest

rates it pays; in this example, all the money would be allocated to weak collateral if the interest-rate differential for weak over strong collateral were 45 basis points or more.

One option the Bank has is to predetermine the spread between rates for strong collateral and weak collateral -- this would correspond to a horizontal line on Figure 2. (Equivalently, the rectangle selected from the feasible set whose upper-right corners are shown in Figure 1b would be constrained to be a fixed number of basis points wider than it was high.) Or, for example, the Bank could fix the proportions of the two kinds of collateral to be used, by choosing a vertical line on Figure 2. Choosing an upward sloping schedule corresponds to "an average of" these approaches.[19]

The point of intersection between the Bank's supply curve and the "demand curve" constructed from the bids shows the interest-rate differential, and the percentage of weak collateral allocated that is generated by the bids in the auction. In this example, it is an interest-rate difference of 27 basis points and a share of weak collateral of 45% -- corresponding to accepting the bids outside the rectangle shown in Figure 1a at the interest rates (5.92% for weak collateral, and 5.65% for strong collateral) corresponding to the rectangle's "northeast" corner.[20]

[19] 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 the assets that were offered at prices furthest below the reference prices; this corresponds to choosing a horizontal supply schedule.

Both suggestions were unduly rigid. The first approach fails to bring different assets into competition with each other, and does not allow the government to move between assets to obtain better bargains. The second approach resolves these problems, but results in the government purchasing large quantities of any asset whose reference price is set too high -- and some mistakes are inevitable, since the government has so much less information than the selling institutions.

Choosing an upward sloping supply schedule maintains the advantage of the reference price approach, while limiting the costs of mispricing them. (This benefit is, of course, additional to the benefits provided by allowing bidders the freedom to use paired bids, etc.)

[20] Note that by specifying the percentage of weak collateral, the intersection of the curves in Figure 2 specifies how any bids that are on the borders of the rectangle, or are on the 45° line, should be treated. In this example, the bidder who bid precisely the cut-off rate for weak collateral for £470 million is allocated £20 million (so that the total allocated on weak collateral is 45% of £2.5 billion equals £1.125 billion); the bidder who bid exactly the cut-off rate for strong collateral for £680 million is allocated £405 million (so that the total allocated against strong collateral is 55% of £2.5 billion equals £1.375 billion).

If the Bank predetermines its supply curve, it can choose whether or not to preannounce it.[21]

Although the example discussed is very simple, the design generalises easily:

More than two classes of collateral

Clearly, there could be more than two classes of collateral, with a cut-off rate for each class, and bids rejected only if they were below *all* of the cut-off rates.

Bidders who only want to offer one kind (or a few kinds) of collateral

As described above, a bidder can refuse to borrow against one or more types of collateral by offering an interest rate of zero for those types.

Bidders who are keen to guarantee winning some given amount

It is also easy for a bidder to guarantee winning a fixed amount by making a bid for this amount of the maximum possible interest rate (or X%, for some arbitrarily large X) against its least preferred collateral, and appropriate discounts (reductions below X%) against the other collaterals.

[In addition to allowing bidders to specify either/or bids for given amounts of money, I recommend explicitly giving them the options of (i) making bids to guarantee winning fixed amounts together with price differential for each such bid, and (ii) making bids against only one kind of collateral -- to draw attention to the possibility of such bids, and

[21] Whether or not it preannounces its supply curve, the Bank may wish to maintain the reserve power to alter it. This mitigates any market power - see the large literature on eliminating "collusive" equilibria in uniform-price auctions. (Klemperer and Meyer, 1989, show random supply reduces the number of equilibria, and Kremer and Nyborg, 2004a, show that with common values and complete information, random supply and discrete bids eliminate low-price equilibria, because discrete bids generate rationing at the equilibrium price, so a small price increase implies a large quantity increase; Back and Zender, 2001, and McAdams, 2007, show the seller can eliminate low-price equilibria by retaining the flexibility to adjust the total quantity after receiving the bids; LiCalzi and Pavan, 2005, discuss the merits of elastic supply; see also Kremer and Nyborg, 2004b, and Kastl, 2008.)

Having the ability to alter the supply curve, therefore, also makes it less likely that the Bank will in fact wish to alter it ex-post.

to aid comprehension. In particular, this emphasises that offering either/or bids is an optional extra that bidders need not take up -- they can restrict themselves to traditional bids against just one kind of collateral if they wish.]

4. Easy Extensions

For our Central Banking application, the features described above may suffice. However, a number of extensions are extremely easy, even in the current graphical framework, and may be more important in other applications. For example, the U.S. TARP might have particularly benefited from the last two extensions in this section -- quantity constraints on groups of bids, and combination bids that guarantee winning minimum amounts, but win more if prices are favourable.

Bidders who want to pay back, or reduce, a previous loan from the Bank

It is easy to include counterparties who may, depending on interest rates, wish to return money to the Bank. Simply add the amount of the loan a counterparty might wish to return, to the total that the Bank makes available to lend, but allow the counterparty to participate in the auction as usual. If the counterparty wins nothing in the auction, the Bank has in effect sold the counterparty's loan on, on the counterparty's behalf. If the counterparty wins the same amount of loan back against the same collateral, there is no change in its position.

"Swappers" who want to change the collateral on a previous loan from the Bank

Exactly the same approach permits a counterparty who wishes to swap the collateral collaterals it is currently using. If, letting the Bank re-auction its current loan, a counterparty in the auction wins the same amount of loan back against the other collateral, it has simply swapped the collateral it is using, and correspondingly changed its interest rate.

Unequal paired bids

One possible extension is to permit either/or bids in which the amounts asked for are different according to the collateral that is offered. For example a bidder might choose as one of its bids "£70 million against ordinary collateral at 3% OR £100 million against weak collateral at 5%". It is not clear that is particularly natural in our Central Banking context,[22] but permitting such bids makes the graphical analysis only slightly more difficult.[23] As usual, this generalises easily to n-dimensions.[24]

Quantity constraints on a group of bids

A bidder may want to specify a bid of the form "£70 million against ordinary collateral at 3% AND £100 million against weak collateral at 5% BUT no more than £140 million in all". Again this is fairly straightforward: on any 45° line on our graph for which the two interest rates are not too different, there is a point which represents a pair of positive cut-off rates that allocates exactly the right amount of the Bank's money in total. Again this can be done in many dimensions, or we can allow bids of the form "X_1 million against C_1 at I_1 % AND …. AND X_n million against C_n at I_n% BUT take no more than m of these n offers", or even bids of the kind "X_1 million against C_1 at I_1% AND …. AND X_n million against C_n at I_n% BUT take only offers satisfying α_1(amount against C_1)+…+ α_n(amount against C_n)$\leq\beta$" for positive constants $\alpha_1,\ldots,\alpha_n,\beta$.

[22] For example, it might be more natural -- and is already permitted by our simpler proposal -- for a bank to make two separate bids (i) "£70 million against ordinary collateral at 3% OR £70 million against weak collateral at 5%" and (ii) "£30 million against weak collateral at 5%". (This would make sense if the bank were interested in a full £100 million, but only had £70 million of strong collateral.)

[23] Assuming a bidder would prefer to be allocated to the collateral that would give it greatest total surplus if it were indifferent about winning at the prices stated in its bid, the bidder's indifference line is no longer a 45° line. (It is indifferent between £70 million against ordinary at 3% and £100 million against weak at 5%; between £70 million against ordinary at 2% and £100 million against weak at 4.3% -- which outcomes give the same *gross* surplus, since (3%-2%)x£70 million=(5%-4.3%)x£100 million=£0.7 million; between £70 million against ordinary at 1% and £100 million against weak at 3.6%; etc.) Instead, as we move inwards (south-west) on any 45° line, we jump at some point from allocating the bidder the smaller amount (£70 million), to allocating the bidder the larger amount (£100 million). It therefore remains true that as we move southwest along any 45° line the amount of funds that the Bank allocates is weakly increasing. So on any 45° line on our graph for which the difference between the two interest rates is not too large, there is a point which represents a pair of positive cut-off rates that allocates exactly the right amount of the Bank's money in total. (If desired, it could alternatively be assumed that a bidder would prefer to be allocated to the collateral that would give it the greatest difference in interest rate relative to the prices stated in its bid.)

[24] For any possible (n-1)-dimensional set of cut-off differences, there is an interest-rate cut-off for the strongest collateral (or any other collateral) that allocates the correct total funds.

More complex constraints

More complex constraints connecting the amounts bid for against the two kinds of collateral can be developed by using several bids in combination.

For example, a counterparty might have enough weak collateral to bid for £100 million, and enough strong collateral to bid for an additional £80 million. The counterparty would like to loan funds against all this collateral if the interest rate against strong collateral is less than 5%, and against weak collateral is less than 7%. However, even if money is expensive, the counterparty would like to borrow an absolute minimum of £40 million.

This can be implemented by making (all of) the following four bids:

1. £40 million at maximum rate against weak OR at maximum rate *less* 2% (=7%-5%) against strong.
2. £80 million at 5% against strong.[25]
3. £100 million at 7% against weak.[26]
4. *minus* £40 million at 7% against weak OR at 5% against strong.

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 cancels the first bid for £40 million “at all costs”.[27]

Using such negative bids, and other extensions, allows us to build more complex constraint structures (including in multiple dimensions), but there are limits to the constraints that can be implemented -- we discuss these in the next section.

[25] Formally, this bid includes “OR at 0% against weak”.
[26] Formally, this bid includes “OR at 0% against strong”.
[27] Clearly not all “negative bids” can be permitted; the negative bid here can be used because the magnitude of the amounts demanded in the second and third bids both exceed the magnitude of the amounts offered up in the negative bid. The fact the cancellation is exact depends upon our uniform pricing rules -- if discriminatory pricing were used, I would recommend slightly amending the rules so that the cancellation becomes exact; see note near end of Appendix 2.

5. 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 standard simultaneous multiple round auction (SMRA).

Begin by considering the case in which the Bank has predetermined the quantity it wishes to offer of each kind of collateral, for example, 7 billion against strong collateral and 3 billion against weak collateral, and the bids represent bidders' true preferences.[28] Then in a standard SMRA, if each bidder bids "myopically" at every step of the SMRA to maximise its profits given these preferences, the outcome will be exactly the one that our procedure selects (in the limit as the bid increments are zero).

That is, imagine that the bidders take turns to make bids in many ascending auctions that are run simultaneously (in this example, 7 billion auctions for a single pound against strong collateral, and 3 billion auctions for a single pound against weak collateral). When it is a bidder's turn, it may make any new bids it wishes, but in doing so must beat any existing winning bid of a competitor by some small bidding increment (say one basis point). (It is not allowed to withdraw any of its own existing bids -- though it would in any case not wish to do so, given the kinds of bids we have discussed.) 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 wishes to submit any new bids. Then the outcome will be exactly the one that our procedure selects (in the limit as the bid increments are zero).

In other words, when bidders behave competitively as price takers, both our mechanism and the SRMA (in the limit as the bid increments are zero) simply select the competitive-

[28] For example, if a bidder's most north-easterly bid is "£100 million at 7 percent against weak OR £100 million at 4 percent against strong" then the bidder is indifferent between any of the outcomes (i) receiving no funds, (ii) receiving £100 million at 7 percent against weak collateral, and (iii) receiving £100 million at 4 percent against strong collateral; for more south-westerly bids, the bidder is indifferent between receiving no funds and receiving a specified collateral *conditional* on having received the same collateral in more north-easterly bids. (If a given bidder's preference between collaterals depends only upon the interest rate difference, all its bids will lie on a 45° line, and it will receive the same collateral on all its bids. If a bidder's preferences are more complicated, it may win funds against different collaterals for different bids, in which case things are more complicated.)

equilibrium price vector (and when this price vector is not unique, they both select the unique vector among these that is lowest in every element).[29]

The case in which the Bank offers a supply curve relating the share of funds allocated to strong collateral to the difference in interest rates, rather than fixing the quantity allocated to each collateral, is not much different. The Bank can be thought of as bidding in the same auction for negative amounts.

It is probably easier, however, to think of the Bank as acting both as the bid-taker selling the maximum possible quantity of both collaterals, and as an additional buyer bidding to buy funds back to achieve a point on its supply curve. That is, in the example in which the Bank wishes to supply £10 billion of funds to the bidders, we consider an SMRA which supplies £20 billion in all -- £10 billion against strong collateral *and* £10 billion against weak collateral; we think of the Bank as an additional bidder with inelastic total demand for £10 billion in funds; and we think of the Bank as bidding for the £10 billion in funds that it demands in exactly the same way as any other bidder in the SMRA. (That is, whenever it is the Bank's turn to bid, it will bid on one or both of the collaterals to both restore its quantity of winning bids to £10 billion and win the quantity against each collateral that puts it back on its supply curve, given the interest-rate differential between the collaterals that it faces.[30])

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

---

[29] To understand this point more generally, 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. In our context, any set of prices such that each bidder takes its preferred allocation is a competitive equilibrium (and the preferences represented by the types of bids we have discussed are consistent with substitutes demand).

[30] The interest rates on winning bids on any collateral will differ by most one bid increment. Ignoring this tiny difference, the Bank will always be able to return to its supply curve, since the interest rate difference between the collaterals can only have increased (decreased) since the Bank's previous turn to bid if the Bank's bids on weak (strong) collateral have all been topped, so the Bank can decrease its quantity of weak (strong) collateral relative to its previous turn to bid, as it will wish to do in this case. (We assume the Bank's supply curve also exhibits substitutes preferences. In this example, that simply means that the quantity supplied against weak collateral is increasing in the interest rate difference -- equivalently, the quantity repurchased against weak collateral is decreasing in the interest rate difference.)

So our auction procedure is equivalent to one in which bidders submit their preferences, and the Bank 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 Bank's supply function may have obscured this, our procedure is symmetric between buyers and sellers.[31])

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, therefore, we could extend our procedure simply by running an SMRA allowing bidders to specify *any* proxy bidding rule -- subject to computational issues. These issues are not very challenging in our simple example,[32] (and nor were they very challenging in the Bank of England's problem).

However, some bidding rules would 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 against collaterals of types A and B $\leq \alpha$) and (total quantity against collaterals of types B and C $\leq \beta$). Then an increase in the interest rate for A might cause the bidder to want to bid for more B instead of A, and this could make it want to bid for less C, so it might regret an earlier offer to bid for a given quantity of C at the current interest rates -- and indeed it might never have made such a bid if the order of play had meant that interest rates evolved differently.[33]

So some constraints on bidding rules are probably desirable. I considered a number of possible extensions to allow more complex forms of bidding than those discussed in the previous sections -- and thinking about what simple preferences a bidder may wish to

[31] The implications of the facts that the Bank is a single player, and has different objectives than its counterparties, are discussed below in Appendix 1.
[32] Obviously, although we described an auction for £10 billion in funds against two possible collaterals as being an auction for 10 billion separate units, this remains a two-dimensional problem since there are just two types of commodity.
[33] The problem is the standard one that competitive equilibrium might not exist without "substitutes preferences" (see Milgrom 2000, Theorems 3, 4). In independent work, Milgrom (2008) explores how to restrict bidders' expressions of preferences to avoid such problems.

express in a SMRA suggests possibilities -- but I did not pursue these very far because it became clear that Bank of England officials favoured a simpler rather than a more complex formulation.

6. Comparison with Alternative Procedures

Comparison with the SMRA

The main reason for proposing our design to the Bank of England was simply that a simultaneous multiple round auction (SMRA) was infeasible because of the time that would be required to run it.

Nevertheless, the design has some additional advantages over the SMRA. It makes exercising market power much harder. Counterparties have an incentive to demand incremental funds at interest rates that are below their valuation (i.e., “unilaterally reduce demand”) in any procedure that charges the same interest rate for all funds against a particular type of collateral,[34] but they are much less likely to have the knowledge to effectively (ab)use their market power in our mechanism because they have to submit their "proxies" before seeing any information about other bidders' demands.

Implicit collusion (coordinated demand reduction) and predatory strategies are also much harder than the equivalent strategies in an SMRA. Not only can the SMRA reveal to bidders where demand reduction or predation may be profitable but -- even more important -- the fact that the SMRA is a dynamic procedure means a bidder can make bids that signal threats and offers to competitors, and also easily punish competitors who fail to cooperate with it.[35, 36]

[34] A myopic proxy bidder does not take into account the impact of its demand for an incremental unit on the price it will end up paying for its inframarginal units, but a bidder can still exercise market power by submitting a proxy consisting of several bids, with all but the first bid at prices below its true value.
[35] Implicit collusion is not impossible in our mechanism, just as it might arise in an ordinary static (sealed-bid) uniform-price auction where a small numbers of competitors with excellent mutual understanding face low uncertainty. (It is more likely if the auction is frequently repeated, and entry of new competitors is hard. For example, implicitly collusive “hockey-stick” bidding has been alleged in electricity-market

The parallel with standard sealed-bid auctions makes our mechanism seem 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.

On the other hand, our procedure limits the kind of strategies that bidders can follow. And even where richer strategies are permitted and feasible, creating the set of constraints that implements them can be complex. In more complicated settings than ours, where it is desirable to let bidders use more sophisticated strategies, an SMRA is a very transparent, and perhaps a more comprehensible, process.

Comparison with running separate static auctions for separate collaterals

Since the timing issues meant a static mechanism was required, a more relevant comparison is between our design and the Bank simply running completely separate static (sealed-bid) auctions, one for strong collateral, and one restricted to weak collateral. (For specificity, we begin by considering ***simultaneous*** static auctions -- we discuss other possibilities below.)

Relative to this alternative, our proposal created flexibility for the bidders (by permitting "paired bids"), and for the bid-taker (by permitting the use of a supply schedule).

Although paired bids superficially seem to make things more complex, they actually simplify the bidding process for all counterparties with a genuine choice of which type of collateral to use: counterparties who use paired bids know they will be allocated money against whichever type of collateral turns out to be more favourable for them.[37] Some bidders might also prefer to make paired bids to avoid giving a bad signal to the Bank by

---

auctions; see Klemperer (2004, 2008).) However, implicit collusion should be even harder to sustain in our mechanism.

[36] Even without deliberately predatory strategies, an SMRA may discourage entry because it gives an advantage to more sophisticated bidders who can better use the information learned between rounds.

[37] That is, at a lower rate relative to their bid.

Paired bids would, of course, anyway always be optional.

As usual, we assume the uniform-pricing rule; with discriminatory auctions things are more complicated (see Appendix 2).

offering weak collateral only. Paired bids minimize the informational disadvantage of bidders who are less well informed about the market (e.g., smaller banks or new bidders).[38]

The ability to choose a supply schedule similarly simplifies the Bank's decision problem, by allowing it to determine the quantity of funds to offer against each collateral as a function of counterparties' actual bids.

Absent the use of paired bids and a supply schedule, auction outcomes will be erratic if counterparties have difficulty judging the rates at which to bid in the two parts of the auction. Most importantly, counterparties' bidding problems might create inefficiency in the form of not getting liquidity to those who value it most. Furthermore (though less important) bidders may be inefficiently allocated to the wrong collateral -- that is, there would be unexploited gains from trade between counterparties.

By forcing bids against different collaterals into competition with each other, the Bank's supply schedule reduces the auction's sensitivity to market power (both monopolistic and oligopolistic), and to manipulation (although that may not be a significant problem in a Central Banking context). Paired bids also help in this respect, by each automatically generating a bid on the "other" collateral (the collateral that would not otherwise have been offered in the given bid). Even if the paired-bid facility is never actually used, its very existence discourages the exercise of market power.

Paired bids increase volume and thicken markets not only because each paired bid automatically creates bids against multiple collaterals, but also by making counterparties more willing to participate in the auction. If there are multiple bidders on both sides of the market, there might be significant additional "second-round" benefits, because the better matching and improved volume and reduced market power on one side of the market would further increase participation and reduce market power on the other side of

[38] The Bank might be concerned that if needier banks were not guaranteed to get money in the auctions they would be more at risk of exploitation by the arbitrageurs.

the market, which would yet further improve volume and reduce market power on the first side of the market, and so on.[39]

In sum, outcomes are more efficient, given the bids that are made; bidders' market power is reduced by the bids for different assets being in competition with each other; and the volumes transacted are greater.

All these things increase overall efficiency, and they all also improve the quality of the information generated by the auction.

They also all yield higher interest rates (i.e., more revenue) for the Bank. The increased volume is obviously beneficial. Reducing market power raises interest rates by reducing incentives for both unilateral (monopsonistic) demand reduction, and "implicitly collusive" (oligopsonistic) demand reduction.[40] Finally, even absent any market power or volume effects, the more efficient allocation of funds generally increases the expected "marginal revenues" of the winning bidders, and hence also raises the interest rates the Bank receives.[41]

Running two static auctions, one for strong collateral, and one restricted to weak collateral, ***sequentially***, would be a slight improvement on running them simultaneously. All parties could adjust their behaviour in the second auction according to the results of the first. Frequently alternating smaller auctions of each type would further increase participants' flexibility, and allow less-informed bidders to wait to learn likely market prices from the early behaviour of more-informed bidders. But the randomness, and thin markets, created by the smaller volumes in each auction would counteract these benefits -- and participants would still have much less flexibility than under our design. Less-frequent alternation of larger-volume auctions (the approach used in the U.K. starting

[39] "Second-round" benefits may be unimportant for a Central Bank which is both non-strategic and the only (or almost only) participant on its side of the market.

[40] See Klemperer (2004, 2008). With fewer bidders, demand-reduction equilibria are usually easier to coordinate on (though there are, at least in theory, cases with discrete "lumpy" units -- not the Central Banking case -- in which these equilibria can arise only when the auctions are joined).

[41] Simply dividing the participants of a single auction among two separate auctions typically lowers the average valuations of the winning bidders, and lower-valuation bidders usually have lower marginal revenues (see Bulow and Roberts (1989), and Bulow and Klemperer (1996)), so separating auctions usually lowers revenues (see Klemperer (2006) and Ellison, Fudenberg and Möbius (2004) for more intuition).

December 2007) gives only limited ability to respond to previous auction results, and also decreases flexibility by limiting the kind of collateral available to bidders at any given time. Thus all these options have similar disadvantages to those of running simultaneous single-collateral auctions, though possibly somewhat smaller in magnitude.[42]

7. 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 because of the time required to run an SMRA. 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 running separate (static) auctions for separate goods, our approach yields:

- *better "matching" between suppliers and demanders*

since separate auctions create randomness because participants do not know how much to attempt to transact in each;

- *reduced market power*

because all bids and offers are forced into competition with each other (and the effects of asymmetries between participants are probably reduced); and

- *greater volume and liquidity in the market*

because the two previous effects increase willingness to participate, and because a "paired" bid automatically generates an additional bid.

[42] Our design is also usually superior to running a single auction that treats all goods as identical, because it achieves a more efficient allocation (and hence also usually higher revenues – see previous note) than treating all goods identically would. One possible exception arises when failing to distinguish between different goods effectively subsidises a weaker bidder, potentially raising revenue (though reducing efficiency). Another exception is when distinguishing too finely between goods increases transactions or information-collection costs (perhaps creating a "cherry picking" prisoner's dilemma in which every participant spends a lot of time working out preferences between close-to-identical products). These exceptions seem more likely to arise in other contexts than in our Central Banking example.

These effects imply our design also generates higher

- *efficiency*
- *revenue* (equivalently, higher interest rates for the Central Bank), and
- *quality of information.*

References

Armantier, Olivier, Sandra Krieger, and James McAndrews (2008). "The Federal Reserve's Term Auction Facility." *Current Issues in Economics and Finance*, Federal Reserve Bank of New York, 14 (5), July

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

Back, Kerry, and Jamie Zender (2001). "Auctions of Divisible Goods With Endogenous Supply." *Economics Letters*, 73, 29–34.

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

Bulow, Jeremy, and Paul Klemperer (1996). "Auctions versus Negotiations." *American Economic Review*, 86, 180-194.

Bulow, Jeremy, and John Roberts (1989). "The Simple Economics of Optimal Auctions." *Journal of Political Economy*, 97, 1060-1090.

Ellison, Glenn, Drew Fudenberg, and Markus Möbius (2004). "Competing Auctions." *Journal of the European Economic Association*, 2, 30-66.

Kastl, Jakub (2008). "Discrete Bids and Empirical Inference in Divisible Good Auctions." Working paper, Stanford University.

Klemperer, Paul (1999). "Auction Theory." *Journal of Economic Surveys*, 13(2): 227-86.

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 (2006). "Answer to Oxford University Economic Theory 2006 Examination." part A4, http://www.nuffield.ox.ac.uk/users/klemperer/OtherTeachingMaterials/2008auctionsolutions.pdf .

Klemperer, Paul (2007). "Bidding Markets." *Journal of Competition Law and Economics*, 3, 1-47.

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

Kremer, Ilan, and Kjell Nyborg (2004a). "Underpricing and Market Power in Uniform Price Auctions." *Review of Financial Studies*, 17, 849-877.

Kremer, IIan, and Kjell Nyborg (2004b). "Divisible Good Auctions - The Role of Allocation Rules." *Rand Journal of Economics*, 35, 147-159.

Krishna, Vijay (2002). *Auction Theory*. New York, NY: Academic Press, US.

LiCalzi, Marco, and Alessandro Pavan (2005). "Tilting the supply schedule to enhance competition in uniform-price auctions." *European Economic Review*, 49, 227-250.

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

Menezes, Flavio M. and Paulo K. Monteiro (2005). *An Introduction to Auction Theory*. Oxford, UK: Oxford University Press.

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

Milgrom, Paul R. (2004). *Putting Auction Theory to Work*. Cambridge, UK: Cambridge University Press.

Milgrom, Paul (2008). "Assignment Messages and Exchanges." Working paper, Stanford University.