

Common-Value Auctions with Liquidity Needs: An Experimental Test of a Troubled Assets Reverse Auction

Lawrence M. Ausubel*, Peter Cramton*, Emel Filiz-Ozbay*,
Nathaniel Higgins†, Erkut Ozbay*, and Andrew Stocking‡

30 May 2011

Abstract

We report the results of an experimental test of alternative auction designs suitable for pricing and removing troubled assets from banks' balance sheets as part of the financial rescue planned by the U.S. Department of Treasury in the fall of 2008. All auction mechanisms tested here are structured so that many individual securities or pools of securities are auctioned simultaneously. Securities that are widely held are purchased in auctions for individual securities; securities with concentrated ownership are purchased as pools of related securities. Each experimental subject represents a bank which has private information about its liquidity need and the true common value of each security. We study bidding behavior and performance of sealed-bid uniform-price auctions and dynamic clock auctions. The clock and sealed-bid auctions resulted in similar prices. However, the clock auctions resulted in substantially higher bank payoffs, since the dynamic auction enabled the banks to better manage their liquidity needs. The clock auctions also reduced bidder error. The experiments demonstrated the feasibility of quickly implementing simple and effective auction designs to help resolve the crisis.

(JEL D44, C92, G01, G21. Keywords: financial crisis, uniform-price auction, clock auction, market design, experiments, troubled assets, TARP.)

* Economics professors at the University of Maryland.

† Economist at the Economic Research Service, USDA.

‡ Market Design Economist at the Congressional Budget Office.

Correspondence to: ausubel@econ.umd.edu, pcramton@gmail.com, filizozbay@econ.umd.edu, nathanielhiggins@gmail.com, ozbay@econ.umd.edu, and astocking@cbo.gov. We thank Power Auctions LLC and its employees for customizing the auction software and making it available for this purpose. The analysis and conclusions expressed in this paper are those of the authors and should not be interpreted as those of the Congressional Budget Office, the Economic Research Service, or the US Department of Agriculture.

1 Introduction

In the fall of 2008, U.S. housing and financial markets were in the midst of severe adjustments. House prices were falling rapidly, and they were expected to continue to fall. Problems in the housing and mortgage markets had spread to a broader array of financial markets. The nation was facing a serious disruption to the functioning of its financial markets that could substantially impair economic activity.¹ The adjustment began following the housing boom that ran from 2003 to early 2006, when delinquencies and foreclosures on mortgages rose, particularly on subprime adjustable-rate mortgage loans (ARMs). Delinquencies also arose for prime ARMs and on so-called alt-A mortgage loans, which were made on the basis of little or no documentation of the borrower's income. Because most mortgages were resold as mortgage-backed securities (MBSs), the rise in delinquencies caused the value of MBSs to decline, in some cases quite sharply.

The problems in mortgage markets spread to the wider financial markets for several reasons. The number of bad mortgages and, consequently, losses on MBSs were expected to be large. The use of complex instruments to fund subprime lending, such as collateralized debt obligations (CDOs), also made it difficult for participants in financial markets to identify the magnitude of the exposure of other participants to losses. Moreover, a number of financial institutions borrowed heavily to finance their mortgage holdings, further increasing their risk exposure. Losses on mortgage assets, and the resulting contraction of the availability of credit to businesses and households, posed a significant threat to the pace of economic activity.

The U.S. Department of Treasury and the Federal Reserve, led by Henry Paulson and Ben Bernanke respectively, considered a host of policy responses to address the illiquidity triggered by market panic and the potential insolvency of many financial institutions. On 3 October 2008, the US Congress passed and the President signed the *Emergency Economic Stabilization Act of 2008* (Public Law 110-343). The Act established the \$700 billion Troubled Asset Relief Program (TARP) which authorized the Secretary of the Treasury to purchase, hold, and sell a wide variety of financial instruments, particularly those that are based on or related to residential or commercial mortgages issued prior to September 17, 2008. The authority to enter into agreements to purchase such financial instruments, which the proposal refers to as troubled assets, would expire two years after its enactment.

An immediate question was what auction designs were well suited to the task. Phillip Swagel, who served as assistant secretary for economic policy at Treasury from December 2006 to 20 January 2009, recalls that in September, "...we were already working hard to set up reverse auctions with which to buy structured financial products such as [mortgage-backed securities], focusing on mechanisms to elicit market prices. On this we received a huge amount of help from auction experts in academia—an outpouring of support that to us represented the economics profession at its finest." (Swagel, 2009, p. 47-48). In a reverse auction, sellers compete with each other to sell a product to a single buyer.

Several potential mechanisms were suggested by market design economists. Ausubel and Cramton (2008a,b) suggested the use of a simultaneous descending clock auction (with some particular features we describe below). Klemperer (2009) suggested a novel sealed-bid auction he dubbed the "product-mix" auction. Treasury settled on a mixed approach: again according to Swagel, "We would have tried two auction approaches, one static and one dynamic—the latter approach is discussed by Lawrence Ausubel and Peter Cramton [2008a], who were among the academic experts providing enormous help

¹ For an overview of the world financial crisis, see French et. al. (2010). For a review of the conditions surrounding the credit markets specifically, see Mizen (2008) and CBO (2008). For a careful study of the conditions in the CDO market that contributed substantially to the crisis, see Barnett-Hart (2009). Gorton (2008,2009) and Swigel (2009) are excellent reviews.

to the Treasury in developing the reverse auctions.” (Swagel, 2009, p. 56) Regardless of the approach used, the Treasury had decided to use “reference prices” in order to purchase many different securities (i.e. securities with many different CUSIPs) in a single auction.²

The use of a reference price is necessary to hold a single auction during which bidders compete to sell a diverse mix of securities. Because ownership of many of the assets was highly concentrated (i.e. competition to sell a single CUSIP would have been relatively low) assets would be grouped together in a pooled-security reverse auction. Each asset is “scored” with a reference price so that the different asset-types can be compared on a single dimension and a single clearing price is determined in the auction.

Reference prices were to be based on Treasury’s best estimates—albeit imperfect estimates—of the value of each CUSIP. Treasury, concerned that poor estimates would be taken advantage of by bidders, considered an auction format in which the reference prices would not be announced until after the auction. Armantier, Holt, and Plott (2010) conducted an experimental test of auctions where the reference price is announced only after bidding has taken place, finding that keeping reference prices secret reduced efficiency and did not save the government money.

The research experiments described here were implemented in October 2008 and were designed to further test several auction mechanisms and design features considered for use by the Treasury. Both experiments include a comparison of sealed-bid and dynamic auctions for many assets (many securities with unique CUSIPs). Experiment 1 tests an auction appropriate for conditions in which ownership of the assets at auction are evenly distributed among banks. Experiment 2 tests an auction mechanism appropriate for conditions in which ownership is instead concentrated unevenly. The auctions in Experiment 2 are very similar to the auctions that would have been used by the Treasury to purchase toxic assets.

Several conclusions emerged from the experiments.

- The auctions were competitive. Owing to the bidders’ liquidity needs, Treasury paid less than the true common value of the securities under either format.
- The sealed-bid auction was more prone to bidder error.
- The dynamic clock auction enabled bidders to manage their liquidity needs better.
- The bidders attained higher payoffs (trading profits plus liquidity bonus) in the dynamic clock auctions than in the sealed-bid auctions.
- Nevertheless, the clock auctions resulted in equivalent aggregate expenditures, so that the benefit to the bidders did not come at the taxpayers’ expense.
- The prices resulting from the clock auctions were a better indication of true values than those from the sealed-bid auctions. We conclude from this that in the context of a troubled-asset crisis like the one facing the Treasury in 2008 the clock auction is apt to reduce risk for both banks and the Treasury, and to generate price information that may help to unfreeze secondary markets.

We conclude that the dynamic clock auction is more beneficial than the sealed-bid auction for both the banks and the taxpayers. The banks attain higher payoffs than in the sealed-bid auction, resulting from better liquidity management. The taxpayers are also better off, as the asset purchase program is

² The acronym “CUSIP” refers to the Committee on Uniform Security Identification Procedures. Each unique security has its own unique CUSIP number, or simply CUSIP, for short.

better directed toward the liquidity needs of the banking sector without increasing the cost of the asset purchase program. The variability of outcomes is also reduced and the informativeness of prices is also increased with the clock format.

The experiment allowed us to do more than compare static and dynamic auctions. More broadly, the experimental format allowed us to create a market design test bed. The test bed helped us to do three things important for all applied market design: (1) demonstrate the feasibility for quick implementation of the auction design; (2) subject the auction design to testing for vulnerabilities; and (3) predict strategic behavior. The commercial auction platform was customized to handle both formats in one week, demonstrating feasibility. Both formats are easy to explain to bidders. Sophisticated subjects required only a three-hour training session to understand the setting, the auction rules, and to practice using the software. Since the auction design was novel and had not been field-tested, a laboratory test was an important part of due diligence. Without the laboratory test bed, we would not have discovered the special vulnerability of the sealed-bid auction to bidder errors. Finally, the test bed was useful in helping to elicit probable strategies from bidders. Again, because the auction format was novel and further because the auction was too complex for equilibrium analysis, bidder behavior could not be predicted without a test bed experiment.

Ultimately, on 12 November 2008, the Treasury decided to concentrate on negotiated equity purchases and postpone the purchase of troubled assets via auction.³ In March 2009, the Treasury proposed auctions to purchase pools of legacy loans from banks' balance sheets, but this time using a forward auction in which private investors compete to buy the pools of loans. Ausubel and Cramton (2009) describe the auction design issues in this new setting and argue for a two-sided auction in which the private investors compete to buy loan pools in a forward auction, and then banks compete in a reverse auction to determine which trades transact. The results we present here are fully applicable to the legacy-loan setting as well. The forward auction is analogous to the security-by-security auction, and the reverse auction is analogous to the reference price auction.

The remainder of the paper is organized as follows. Section 2 we summarize the experimental literature with respect to dynamic and sealed-bid auctions. Our analysis builds on this literature. Section 3 briefly describes the experimental setup. The instructions and related materials are available in the appendix. Section 4 provides an econometric analysis of the results. Section 5 describes the implementation and results of a recombinant procedure, a procedure that explores the full range of outcomes in the sealed-bid auction. Section 6 concludes.

2 Dynamic vs. static auction designs

There is a rich economic literature that points to the advantages of a competitive process over negotiation (see e.g., Bulow and Klemperer, 1996) and thus we focus exclusively on the use of auctions to accomplish Treasury's objectives. It is within this context that we designed our auction experiment to help us understand the outcomes and relative advantages of alternative auction formats. One of the initial decisions facing Treasury was whether to conduct a static (sealed-bid) or dynamic (descending-bid) auction.⁴

A frequent motivation for the use of dynamic auctions is reducing common-value uncertainty (Milgrom and Weber 1982). In the troubled-asset setting there is a strong common-value element: A

³ See http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aoD4Qeblb_Fo&refer=home

⁴ See Ausubel and Cramton (2004, 2006), Cramton (1998), McAfee and McMillan (1987), and Milgrom (2004) for further discussion.

security's value is closely related to its "hold to maturity value," which is roughly the same for each bidder. Each bidder has an estimate of this value, but the true value is unknown. The dynamic auction, by revealing market supply as the price declines, lets the bidders condition their bids on the aggregate market information. As a result, common-value uncertainty is reduced and bidders will be comfortable bidding more aggressively without falling prey to the winner's curse—the tendency in a procurement setting of naïve sellers to sell at prices below true value.

In the context of many securities, the price discovery of a dynamic auction plays another important role. By seeing tentative price information, bidders are better able to make decisions about the quantity of each good to sell. This is particularly useful because the values of securities are related. Bidding in the absence of price information makes the problem much more difficult for bidders. Furthermore, with a dynamic auction, the bidder is better able to manage both liquidity needs and portfolio risk. In contrast, managing liquidity needs in a simultaneous sealed-bid auction is almost impossible.

Another advantage of a dynamic auction is transparency. Each bidder can see what it needs to do to win a particular quantity. If the bidder sells less, it is the result of the bidder's conscious decision to sell less at such a price. This transparency is a main reason for the high efficiency of the descending clock auction in practice.

Finally, as a practical matter, a clock auction allows for feedback between auction rounds, reducing the likelihood that a mistaken bid will go undetected. Bidders do make mistakes, entering bids incorrectly because of keystroke or other human error. A recent example occurred in the Mexican Central Bank's auction for U.S. currency on 19 May 2009. A bank entered an erroneous bid that caused it to overpay by US\$355,340. All other accepted bids in the auction were within 0.3% of the exchange rate traded that day, while the erroneous bid was 7.4% greater than the concurrent rate.⁵ Reducing the likelihood of bidder error is important. We provide evidence in this paper that bidder error is less likely under the clock format than under the sealed-bid format.

The experimental economics literature strongly supports the conclusion that dynamic auctions outperform sealed-bid auctions in terms of efficiency and price discovery. In sealed-bid auctions there is a tendency to consistently overbid (Kagel, Harstad, and Levin 1987; McCabe, Rassenti, and Smith 1990), often resulting in inefficient outcomes. In contrast, many laboratory and field experiments have demonstrated that the clock auction format is simple enough that even inexperienced bidders can quickly learn to bid optimally (Kagel, Harstad, and Levin 1987). Kagel (1995) finds that bidders readily transfer the experience gained in sealed-bid auctions to the clock auction format. Bidders in Levin, Kagel, and Richard (1996) appear to adopt simple strategies that incorporate dynamically changing information from the clock auction, namely the prices at which other bidders drop out, and efficient outcomes are obtained.

A principal benefit of the clock auction is the inherent price-discovery mechanism that is absent in any sealed-bid auction. Specifically, as the auction progresses, participants learn how the aggregate demand changes with price, which allows bidders to update their own strategies and avoid the winner's curse (Kagel 1995). Levin, Kagel, and Richard (1996) show that bidders suffer from a more severe winner's curse in the sealed-bid format than in a clock auction. Kagel and Levin (2001) compare a clock auction and a sealed-bid auction when bidders demand multiple units, and confirm that outcomes are much closer to optimal in the clock auction. Efficiency in the clock auction always exceeded 97%.

5

See www.banxico.org.mx/eInfoFinanciera/InfOportunaMercadosFin/MercadoCambios/ResultadosSubastas/PosturasAsignadasSubCamConvoc1

Moreover, in the Ausubel auction (a particular type of clock auction, see Ausubel 2004, 2006) bidders achieve optimal outcomes 85.2% of the time, as compared to only 13.6% of the time in a sealed-bid auction. McCabe, Rassenti and Smith (1990) found 100% efficient outcomes in 43 of 44 auctions using a clock auction. Kagel and Levin (2008) provide further evidence of more efficient outcomes with a clock format in the multi-unit setting. Alsemgeest, Noussair, and Olson (1996) also find that clock auctions are efficient both in single and multi-unit supply scenarios, achieving better than 99.5% efficiency and 98% efficiency.

The principle advantage of a sealed-bid auction is its apparent simplicity and relatively “inexpensive” setup. Some would argue that a sealed-bid auction is also less vulnerable to collusion. Some also fear that even a quick dynamic auction would expose participants to significant unhedged positions as a result of real time interactions with financial markets. This latter complaint can be addressed by conducting the auction when the major financial markets are closed.

3 Experimental setup

During the period 12-24 October 2008 and 6-11 November 2008, using commercial auction software customized for our purpose, we tested two different auction environments at the University of Maryland’s experimental economics laboratory.⁶ The objective of the experimental setup was to mimic the environment faced by the Department of Treasury. Specifically, Treasury faced the challenge of purchasing assets so as to balance two competing criteria: 1) assuring that the taxpayer would not overpay for the assets; and 2) improving banking sector stability by purchasing assets from those banks most in need of liquidity. Ausubel and Cramton (2008a,b) discussed the design issues as they appeared in October 2008 and proposed a specific auction format. The experiments described below were designed to provide insights into bidding behavior and performance of that format relative to alternative formats, as well as demonstrate the feasibility of quickly implementing either format of auction as part of the financial rescue. For each auction format and information setting analyzed, we compare sealed-bid uniform-price auctions with dynamic clock auctions, varying the level of competition, information, and banks’ need for liquidity.

The experimental auction environments were closely tailored to the likely settings of the planned auctions for troubled assets. Specifically, to model the case where there was sufficient competition to conduct a competitive auction for individual securities, we ran an 8-security simultaneous reverse auction. Each security had a pure common value with unconditional expectation of 50 cents on the dollar, bidders had private information about the common value, and a fixed quantity of each security was purchased in the same reverse auction. This is what we refer to as a *security-by-security* auction. In the second auction environment, the ownership of the security was too concentrated to allow individual purchase. Securities of a similar quality were instead pooled together, thus mitigating the concentration-of-ownership problem. In this second auction environment each security had a pure common value, bidders had asymmetric endowments, and bidders with larger holdings of a security had more private information about the common value. In order to implement an auction where dissimilar items are purchased together, bidders compete on the basis of a reference price, which reflects the government’s best estimate of the security’s value. Bidders then compete on a relative basis—a bid expresses willingness to tender a security at a stated percentage of the security’s reference price. This is what we refer to as a *reference price auction*.

⁶ <http://www.econ.umd.edu/resources/computing/experimental>

The human subjects bidding in the auctions were experienced PhD students, highly motivated by the prospect of earning roughly \$1200 each—the actual amount depending on performance—for participating in twelve experimental sessions, each lasting two to three hours, over the three-week period. We chose to use experienced PhD students for these experiments, since the environment is considerably more complex than a typical economics experiment, and we believed that the PhD students' behavior would be more representative of the sophisticated financial firms who would be participating in the actual auctions.

In terms of scope, the experimental “banks” held roughly 8,000 distinct troubled securities, potentially available for purchase. For the purposes of this paper, these assets fall into two general groups: 1) those securities with ownership concentrated among only a few firms; and 2) those securities with less concentrated ownership. By the nature of these troubled assets, both the banks and the government believed them to be worth less than face value. However, some securities are more “troubled” than others. Some are relatively high-valued securities (e.g., a market value of 75 cents on the dollar) and others are relatively low-valued securities (e.g., a market value of 25 cents on the dollar).

For purposes of exposition we describe the two auction environments as Experiment 1, an 8-security simultaneous reverse auction, and Experiment 2, a pooled security reverse auction. Experiments 1 and 2 were conducted over a total of 12 sessions. The schedule of treatments is given in Table 1. The individual and pooled auctions are described below, with more detail provided in the appendix. Each session involved four auctions in the order indicated.

Table 1. Schedule of treatments

Order of Treatment:	First	Second	Third	Fourth	
Session					
Positive Liquidity Need					
1 - 4	Auction Type	Sealed-bid	Sealed-bid	Clock	Clock
	# of bidders	4	8	4	8
	reference prices	NA	NA	NA	NA
5, 7	Auction Type	Sealed-bid	Clock	Sealed-bid	Clock
	# of bidders	8	8	8	8
	reference prices	More precise	More precise	Less precise	Less precise
6, 8	Auction Type	Sealed-bid	Clock	Sealed-bid	Clock
	# of bidders	8	8	8	8
	reference prices	Less precise	Less precise	More precise	More precise
Zero Liquidity Need					
9 - 10	Auction Type	Sealed-bid	Clock	Sealed-bid	Clock
	# of bidders	4	4	4	4
	reference prices	NA	NA	NA	NA
11 - 12	Auction Type	Sealed-bid	Clock	Sealed-bid	Clock
	# of bidders	8	8	8	8
	reference prices	Less precise	Less precise	Less precise	Less precise

3.1 Experiment 1: 8-security simultaneous reverse auction

In Experiment 1, bidders compete to sell their symmetric holdings of eight securities to the Treasury. Two formats are used:

- *Simultaneous uniform-price sealed-bid auction (“sealed-bid auction”).* Bidders simultaneously submit supply curves for each of the eight securities. Supply curves are non-decreasing (i.e. upward-sloping) step functions. The auctioneer then forms the aggregate supply curve and crosses it with the Treasury’s pre-announced and fixed demand. The clearing price is the lowest-rejected offer. All quantity offered below the clearing price is sold at the clearing price. Quantity offered at the clearing price is rationed to balance supply and demand, using the proportionate rationing rule⁷.
- *Simultaneous descending clock auction (“clock auction”).* The eight securities are auctioned simultaneously over multiple rounds. In each round, there is a price “clock” that indicates the

⁷ The proportionate rationing rule only plays a role in the event that multiple bidders make reductions at the clearing price. The rule then accepts the reductions at the clearing price in proportion to the size of each bidder’s reduction at the clearing price. Thus, if a reduction of 300 is needed to clear the market and two bidders made reductions of 400 and 200 at the clearing price, then the reductions are rationed proportionately: the first is reduced by 200 and the second is reduced by 100. The actual reduction of the first bidder is twice as large as the second bidder, since the first bidder’s reduction as bid is twice as large as the second bidder’s reduction.

start of round price and end of round price per unit of quantity. Bidders express the quantities they wish to supply at prices they select below the start of round price and above the end of round price. At the conclusion of each round, bidders learn the aggregate supply for each security. In subsequent rounds, the price is decremented for each security that has excess supply, and bidders again express the quantities they wish to supply at the new prices. This process repeats until supply is made equal to demand. The tentative prices and assignments then become final. Details of the design are presented in Ausubel and Cramton (2008).

Six sessions were dedicated to Experiment 1 to test the following three auction attributes: 1) the effect of sealed-bid vs. clock formats; 2) the effect of liquidity needs; and 3) the effect of increased competition. In sessions 1-4, we conducted paired sealed-bid and clock auctions with both low and high levels of competition (a total of four bidders competed in low-competition auctions, while eight competed in a high-competition auctions). Sessions 9-10 were similar, except that bidders did not have liquidity needs. That is, subjects were not given a bonus based on the sale of securities during the auction. Instead, a subject's take-home pay was based entirely on the profits they made when they sold a security to the government for more than its true value. We focused on the low-competition case in sessions 9-10, substituting an extra pair of 4-bidder auctions in place of the 8-bidder auctions. As a result of this schedule in sessions 9 and 10, we effectively gave players four auction pairs (sealed-bid and clock) of learning in two consecutive days, focused only on the 4-bidder auction.

The experimental design was intended to facilitate a direct comparison of the sealed-bid auction and the clock auction. Before each sealed-bid auction, each bidder learned the realizations of one or more random variables that were relevant to the value of the securities that she owned. The same realizations of the random variables applied to the clock auction immediately following the sealed-bid. Thus, in successive pairs of experimental auctions, the securities had the same values and the bidders had the same information. Bidders were not provided with any information about the outcome of a given sealed-bid auction before the following clock auction, in order to avoid influencing the behavior in the clock auction.⁸

The value of each security in cents on the dollar is the average of eight iid random variables uniformly distributed between 0 and 100:

$$v_s = \frac{1}{8} \sum_{i=1}^8 u_{is}, \text{ where } u_{is} \sim_{iid} U[0,100],$$

where a bidder's private information about security s is the realization u_{is} . This is true both for the 8-bidder and 4-bidder auctions, so that only the first four draws are revealed in the 4-bidder auction. This design allowed the true values to have the same distribution in both 4-bidder and 8-bidder auctions which caused the private information to have the same precision.

A bidder profits by selling securities to the Treasury at prices above the securities' true values. Profit (in million \$) is defined as:

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_{s=1}^8 (p_s - v_s) q_{is},$$

where the quantity sold is q_s of security s at the price p_s .

⁸ Observe that, inherently, information about a clock auction *must* be revealed, as bidders learn aggregate information about Round 1 before the start of Round 2, etc. Thus, it would have been impossible to run the sealed-bid auctions after the clock auctions without influencing the behavior in the sealed-bid auctions.

In sessions 1-4, bidders also have a need for liquidity. The sale of securities to the Treasury is the source of a bidder's liquidity. The liquidity need, L_i , is drawn iid from the uniform distribution on the interval [250, 750]. Bidders know their own liquidity need, but not that of the other bidders. Bidders receive a bonus of \$1 for every dollar of sales to the Treasury up to their liquidity need of L_i . Bidders do *not* get any bonus for sales to the Treasury above L_i . Thus, their bonus is:

$$\min \left[L_i, \frac{1}{100} \sum_{s=1}^8 p_s q_{is} \right].$$

Given that bidders care about both profits and liquidity, their total payoff is the combination of the two:

$$U_i(p, q_i, v) = \begin{cases} \frac{1}{100} \sum_{s=1}^8 (2p_s - v_s) q_{is} & \text{if } \frac{1}{100} \sum_{s=1}^8 p_s q_{is} < L_i \\ L_i + \frac{1}{100} \sum_{s=1}^8 (p_s - v_s) q_{is} & \text{otherwise} \end{cases}.$$

In each session, two auctions were selected at random (one from each pair of auctions) to determine bidders' take-home earnings. We used a conversion factor of \$1 in take-home pay for every \$10 million in experimental earnings.

Given the relatively tractable theoretical nature of the experimental setup without the liquidity constraint, we calculated a benchmark bid based on equilibrium bidding strategies in a common value auction (Milgrom and Weber, 1982):

$$\text{4-bidder sealed-bid strategy: } b_{is} = \frac{1}{8} \left(2u_{is} + 2 \left(\frac{u_{is} + 100}{2} \right) + 4 \cdot 50 \right) = \frac{3}{8} u_{is} + \frac{75}{2}.$$

$$\text{8-bidder sealed-bid strategy: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + 3 \left(\frac{u_{is} + 100}{2} \right) \right) = \frac{5}{8} u_{is} + \frac{75}{4}.$$

These Bayesian Nash equilibrium strategies are based on a theoretical framework that differs from our experiment in two ways: 1) they ignore any behavioral adjustments resulting from the liquidity bonus; and 2) they assume that bidders sell their holdings to the Treasury as an indivisible block (i.e. either their entire endowment or nothing). Despite that, the benchmark strategies provide guidance in a static or dynamic setting. As a result, these strategies were explained to bidders and made operational in a bidding tool (i.e., the bidding tool facilitated updating of the strategy following a drop in supply by backwardly inducting the values of the bidder who reduced their supply). Assuming all players play the benchmark strategy, we simulated both the sealed-bid and clock auctions under the two competition levels. These simulations provide an expected clearing price for each of the 8 securities as well as bidder-specific profits and payoffs. While in the experimental auction we anticipated that the liquidity bonus would be likely to cause players to bid more aggressively than predicted by the benchmark, this behavioral change was not included in our simulations.

3.2 Experiment 2: pooled security reverse auction

In the reference price auctions, the holdings of the eight individual securities are too concentrated for there to be competitive auctions on a security-by-security basis. Think of a reverse auction for apples and oranges. In a simultaneous auction all bidders would submit bids to sell apples and oranges at the same time—apple-bids would compete against other apple bids and orange-bids would compete against other orange-bids. The result of the auction would be two separate clearing prices, one for apples and

one for oranges. In a pooled-fruit auction apple-bids would compete against other apple-bids *and* all the orange-bids. Since apples and oranges are clearly different fruits, in order to consider the relative merit of apple-bids and orange-bids the auctioneer would state a price that she believes to be a fair price for apples (say \$0.50), as well as a price that she believes to be a fair price for oranges (say \$0.75). Bids are then ranked according to the discount on the estimated value, so that an apple-bid of \$0.50 (a discount ratio of 1) would rank as more *expensive* than an orange-bid of \$0.60 (a discount ratio of 0.8).

The defining features of the pooled auction are as follows:

- The clearing prices for different securities (i.e., securities with different CUSIP numbers) are determined within the same auction;
- Bidder endowment and thus price signals are asymmetric for each security;
- Before an auction, the Treasury determines and announces its estimate of the value of each security—these are referred to as *reference prices*;
- The prices in the sealed-bid auction, or in each round of the descending-clock auction, are expressed as a percentage of the reference price for each security—these are referred to as *price points*; and
- Clearing occurs when the cost of purchasing the securities offered at a given price point equals the budget allocated for the auction.

As in Experiment 1, two auction formats are considered:

- *Simultaneous uniform-price sealed-bid*. Bidders simultaneously submit supply curves for each of the securities within the pool. Supply curves are upward-sloping step functions, where prices are expressed as price points (a percentage of the reference price) and quantities are expressed in dollars of face value. The auctioneer then forms the aggregate supply curve and equates it with the Treasury's demand. The clearing price is the lowest rejected offer. All securities offered at price points below the clearing price point are purchased at the clearing price point. Securities offered at exactly the clearing price point are rationed by a proportional rationing rule.
- *Simultaneous descending-clock*. There is a price "clock" indicating the current range of price points. For example, in Round 1, bidders express the quantities that they wish to supply of each security at all price points from 106% to 102% of the respective reference price for securities within that auction. After Round 1, the auctioneer aggregates the individual bids and informs bidders of the aggregate quantity that was offered at 102%. Assuming that supply exceeded demand, the price is decremented; for example, in Round 2, bidders may express the quantities that they wish to supply of each security at all price points from 102% to 98% of the respective reference prices. The process is repeated, with the price decremented, bids submitted and quantities aggregated, until supply is made equal to demand. Then, as in the sealed-bid auction, all securities offered at price points below the clearing price point are purchased at the clearing price point, and bids at exactly the clearing price point are rationed by a proportional rationing rule.

Details of the designs are described in Ausubel and Cramton (2008).

Six sessions were dedicated to test the following three auction attributes: 1) the effect of sealed-bid vs. clock auction format; 2) the effect of the liquidity bonus; and 3) the effect of increasing precision with respect to the reference price. In sessions 5-8, we ran a low precision sealed-bid and clock auction

and high precision sealed-bid and clock auction (four auctions total) in that order⁹. Thus bidders completed four auction pairs (sealed-bid and clock) for each of the low precision and high precision auctions (one pair of each precision level each day). In sessions 5 and 6, we removed the liquidity bonus and ran two low precision sealed-bid and clock auctions per session for a total of four auctions in each session. As a result, we effectively gave players four auction pairs (sealed-bid and clock) of learning in two days, but only with the low precision auction.

Table 2. Holdings of securities by bidder and security in million \$ of face value

	High-Quality Securities				Low-Quality Securities				Total
	H1	H2	H3	H4	L1	L2	L3	L4	
Bidder 1	20	0	0	0	0	5	5	10	40
Bidder 2	0	20	0	0	10	0	5	5	40
Bidder 3	0	0	20	0	5	10	0	5	40
Bidder 4	0	0	0	20	5	5	10	0	40
Bidder 5	0	5	5	10	20	0	0	0	40
Bidder 6	10	0	5	5	0	20	0	0	40
Bidder 7	5	10	0	5	0	0	20	0	40
Bidder 8	5	5	10	0	0	0	0	20	40
Total	40	40	40	40	40	40	40	40	
Expected price	75	75	75	75	25	25	25	25	
Expected value	30	30	30	30	10	10	10	10	
Total value	120				40				

Bidder endowments for each security are described in Table 2. Each bidder had an endowment of \$40 million of face value, divided differently across securities. Similarly, there are \$40 million of face value for each security. Treasury has a demand for 25% of the total face value within each pool of securities, which might be involve the purchase of one or more individual securities.

The value of each high-quality security $s \in \{H1,H2,H3,H4\}$ in cents on the dollar is the average of n iid random variables uniformly distributed between 50 and 100:

$$v_s = \frac{1}{n} \sum_{j=1}^n u_{js}, \text{ where } u_{js} \sim_{iid} U[50,100].$$

The value of each low-quality security $s \in \{L1,L2,L3,L4\}$ in cents on the dollar is the average of n iid random variables uniformly distributed between 0 and 50:

$$v_s = \frac{1}{n} \sum_{j=1}^n u_{js}, \text{ where } u_{js} \sim_{iid} U[0,50].$$

⁹ See footnote 5. , In sessions 5 and 7, the more precise sealed bid and clock auctions were conducted first. In sessions 6 and 8, the less precise sealed bid and clock auctions were conducted first.

For auctions with more precise reference prices, $n = 16$; for auctions with less precise reference prices $n = 12$. The reference price r_s for security s is given by

$$r_s = \frac{1}{n-8} \sum_{j=9}^n u_{js}.$$

Thus, the reference price is based on eight realizations in the more precise case (1/2 of all realizations) and on four realizations in the less precise case (1/3 of all realizations). Reference prices are made public before each auction starts.

For each \$5 million of security holdings, bidder i receives as private information one of the realizations u_{js} . Thus, bidder 1, who holds \$20 million of security 1, gets four realizations (see Table 2). In this way, those with larger holdings have more precise information about the security's value. Observe that this specification requires the holders of each given security to receive collectively a total of eight realizations. Since there are eight realizations available (besides the ones that form the reference price), each of the realizations u_{js} ($i = 1, \dots, 8$) can be observed by exactly one bidder.

Suppose that the auction clearing price-point is p_H for the high-quality pool and p_L for the low-quality pool, where the price-point in the auction is stated as a fraction of the reference price. Then $p_s = p_H r_s$ for $s \in \{H1, H2, H3, H4\}$ and $p_s = p_L r_s$ for $s \in \{L1, L2, L3, L4\}$.

If a bidder sells the quantity q_s of the security s at the price p_s , then profit is

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_s (p_s - v_s) q_{is},$$

where the 1/100 factor converts cents into dollars. As with Experiment 1, when bidders have a liquidity need (sessions 5-8), it is drawn iid from the uniform distribution. In Experiment 2, however, the cash scale is increased, and thus liquidity needs are drawn from the interval [2500, 7500]. Each bidder knows his own liquidity need, but not that of the other bidders. The bidder receives a bonus of \$1 for every dollar of sales to the Treasury up to his liquidity need:

$$\min \left[L_i, \frac{1}{100} \sum_s p_s q_{is} \right].$$

Combining the profit and the liquidity penalty results in the bidder's total payoff

$$U_i(p, q_i, v) = \begin{cases} \frac{1}{100} \sum_s (2p_s - v_s) q_{is} & \text{if } \frac{1}{100} \sum_s p_s q_{is} < L_i \\ L_i + \frac{1}{100} \sum_s (p_s - v_s) q_{is} & \text{otherwise} \end{cases}$$

Thus, an additional dollar of cash is worth two dollars when the bidder's liquidity need is not satisfied, but is worth one dollar when the liquidity need is satisfied. To be roughly comparable to Experiment 1, bidder's take-home pay was calculated such that they received \$1 in take-home pay for every \$100,000 in experimental earnings.

Unlike Experiment 1, there is no Bayesian Nash equilibrium bidding strategy for a similar auction that we can use as a benchmark. The reference price auction is beyond current theory. The pooling of securities combined with the use of reference prices violates monotonicity in signals, meaning that a higher signal does not necessarily translate to a higher bid. Monotonicity between signals and value exists within a particular security (i.e., ceteris paribus, a higher signal suggests a higher value); however, there is no monotonicity across securities within a particular pool.

Monotonicity holds when a higher signal implies a higher expected value to the bidder. This relationship is broken by the existence of the reference price. Consider that a security with a higher reference price has a higher expected value to each bidder, all else equal. Holding the common value of a security fixed, bidders prefer a higher reference price, since a high reference price makes the security more competitive in the pool. Thus, in determining her bid, a bidder must consider the countervailing forces of signals and reference prices. It is difficult to recommend how a bidder should respond to a high signal with a low reference price, a low signal with a high reference price, etc.

3.3 Experimental subjects

The training of subjects and all experimental sessions took place in the Experimental Lab of the University of Maryland's Economics Department. This is a new state-of-the-art facility for conducting economic experiments. Each subject has her own private cubical with computer and necessary software. The subject pool consisted of Ph.D. students at the University of Maryland and George Mason University. The students had taken or are taking an advanced graduate course in game theory and auction theory, and are pursuing degrees in economics, business, computer science, or engineering.

In each session of approximately three hours, 16 bidders, out of a total subject pool of 19, participated in four auctions. Each auction consisted of four or eight bidders (i.e., there were always multiple auctions conducted in parallel) and the bidders were randomly and anonymously matched.

Bidders' payoffs consisted of the sum of two terms. First, each bidder received trading profits according to the difference between the common value, v , of the security, and the price, p , at which the bidder's securities were purchased. Hence, if the bidder sold a quantity, q , of securities, the bidder's trading profits equaled: $q \cdot (p - v)$. Second, each bidder was randomly assigned a *liquidity need*, L , and received an additional dollar of payoff for each dollar in sales, $q \cdot p$, up to L that the bidder received in a given auction.

At the conclusion of all sessions, each subject received a check equal to a show-up fee of \$22 per session plus an amount proportional to her total experimental payoff as described above. Average take-home pay was \$100.43 per session.

The next section describes the results.

4 Experimental results

The primary results comparing sealed-bid and clock aspects of the two experiments are summarized in Tables 3-6. First considering just the results from Experiment 1 with the liquidity bonus, we see that even though clearing price and profits are statistically indistinguishable between the two auction formats, the variability of profit is much higher in the sealed-bid auction compared to the clock. Thus the results from the clock auction would appear to be more stable and predictable. Treasury would appear to best satisfy their first objective to consistently get the best possible price for the taxpayers using a clock auction, though the difference is small. This is particularly important when a liquidity bonus is in effect; without the liquidity bonus (Table 4) profits are statistically greater than zero and the clock profits are significantly higher for the clock auction (178) relative to the sealed-bid auction (118) with negligible differences between the standard deviation.

Table 3. Comparison of mean outcomes by auction type in Experiment 1 with liquidity bonus

Variable	Auction Type		Result
	Sealed-Bid	Clock	
Clearing Price	47.79 (1.41)	49.57 (1.32)	The clearing price is statistically indistinguishable for the Clock and Sealed Bid auction (t-test p-value of 0.3621)
Profit	-39.54 (21.1)	-13.03 (15.5)	Profits are statistically indistinguishable between the two auction formats (t-test p-value of 0.3135)
Standard Deviation of Profit	239.3	175.8	Higher standard deviation of profit in sealed-bid than clock (variance ratio test p-value 0.0006)
Liquidity Bonus	428 (16)	466 (13)	Clock liquidity bonus is significantly larger than sealed-bid liquidity bonus (t-test p-value of 0.0562)
Payoff	388 (25)	453 (20)	Clock payoff is significantly higher than sealed-bid payoff (t-test p-value of 0.0400)
Standard Deviation of Payoff	281.7	221.5	Higher standard deviation of payoff in sealed-bid than clock (variance ratio test p-value 0.0071)
Overshooting the liquidity need	692 (41)	605 (37)	Overshooting the liquidity need is almost significantly less in clock than in sealed-bid (t-test p-value of 0.1210)

Note: mean value is shown with standard error in parentheses

Table 4. Comparison of mean outcomes by auction type in Experiment 1 without liquidity bonus

Variable	Auction Type		Result
	Sealed-Bid	Clock	
Clearing Price	55.81 (0.66)	58.82 (0.51)	The clearing price is significantly higher for the Clock auction (t-test p-value of 0.0004)
Profit = Payoff	118.07 (14.89)	178.20 (14.16)	Profits are significantly greater than zero in both cases, and are significantly higher in the Clock auction (t-test p-value of 0.0041)
Standard Deviation of Payoff	119.1	113.3	Standard deviation of payoff in sealed-bid is statistically identical to that of clock (variance ratio test p-value 0.6919)

Note: mean value is shown with standard error in parentheses

Turning to Treasury's second objective related to buying assets from those banks most in need of liquidity, we examine the payoffs from the two auction formats. Payoffs are significantly higher under the clock auction (453) compared to the sealed bid (388). We also see that the variability of total payoffs is higher under the sealed-bid auction than the clock which supports the premise that the additional information provided by the clock auction format leads to more consistent, less variable outcomes. Once again, Treasury is best served in achieving their second objective with a clock auction.

Turning to Experiment 2 with liquidity need, we see that there is no difference in the clearing price between the two auction formats and while the profits are lower in the clock auction (-799) compared to the sealed-bid auction (-693), the difference is not significant. In addition, there is not a significant variation in the standard deviation of the profit. This result is mimicked in Table 6 when the liquidity bonus is not present. In terms of achieving Treasury's first objective, the two auction formats would seem indistinguishable.

Table 5. Comparison of mean outcomes by auction type in Experiment 2 with liquidity bonus

Variable	Auction Type		Result
	Sealed-Bid	Clock	
Clearing Price	85.22 (0.81)	83.87 (1.18)	The clearing pricepoint is significantly indistinguishable between the two auction formats (t-test p-value of 0.3485)
Profit	-693.06 (51)	-798.60 (57)	Profits are significantly less than zero in both cases, but no significant difference in mean profits (t-test p-value of 0.1680)
Standard Deviation of Profit	574.1	645.2	No significant difference in the standard deviation on profit in clock compared to sealed-bid (variance ratio test p-value 0.1896)
Liquidity Bonus	3915.0 (172)	4517.2 (131)	Clock liquidity bonus is significantly larger than sealed-bid liquidity bonus (t-test p-value of 0.0059)
Payoff	3222.0 (146)	3718.6 (116)	Clock payoff is significantly higher than sealed-bid payoff (t-test p-value of 0.0083)
Standard Deviation of Payoff	1653.5	1311.9	Higher standard deviation of payoff in sealed-bid than clock (variance ratio test p-value 0.0095)
Overshooting the liquidity need	1984.0 (290)	904.8 (154)	Overshooting the liquidity need is less in clock than in sealed-bid (t-test p-value of (0.0014)

Note: mean value is shown with standard error in parentheses

Table 6. Comparison of mean outcomes by auction type in Experiment 2 without liquidity bonus

Variable	Auction Type		Result
	Sealed-Bid	Clock	
Clearing Price	93.7 (1.41)	94.6 (1.36)	The clearing pricepoint is significantly indistinguishable between the two auction formats (t-test p-value of 0.6716)
Profit = Payoff	160.1 (39)	244.6 (38)	Profits are significantly greater than zero, and are almost significantly higher in the Clock auction (t-test p-value of 0.1215)
Standard Deviation of Payoff	309.7	303.5	Standard deviation of payoff in sealed-bid is statistically identical to that of clock (variance ratio test p-value 0.8739)

Note: mean value is shown with standard error in parentheses

However, when the liquidity bonus is included in the analysis (Table 5), we see that the mean payoff under the clock auction (3,718) is significantly higher than the payoff under the sealed-bid auction (3,222). Moreover, the standard deviation of payoff is higher under the sealed-bid and the magnitude by which experimental subjects overshoot their liquidity need was higher in the sealed bid (1,984 sealed bid overshoot and 905 clock overshoot). Both of these results suggest that the clock auction is a more efficient and accurate means of helping the Treasury determine which banks are most in need of liquidity and allowing the banks to best manage their need for liquidity.

In the following two sections we discuss in more detail the econometric analysis of the data.

4.1 Experiment 1: simultaneous descending clock

The baseline regression results demonstrating the effect of liquidity and learning across the six sessions of auctions in Experiment 1 are shown in Table 7. There are three striking results from this table. First, we see the results described in Tables 3-6: the profit between the sealed bid and clock auction (regression 2) is statistically identical; whereas, when the liquidity bonus is zero, bidders earn a significantly higher profit in the clock auction (\$60). Theory predicts that without the liquidity bonus, the expected payoff from the clock and sealed bid auctions should be identical, though the sealed bid is likely to have higher profit variance. This is not borne out in the results and may be because additional information made available to bidders in the clock auction facilitated tacit collusion. In the auctions with liquidity, tacit collusion was more complicated to implement due to the multiple bidder objectives.

When the liquidity bonus is included in the payoff, the clock auction generates significantly higher payoffs (\$65) relative to the sealed bid auction.

Table 7. Experiment 1 regression results with experimental subject fixed effects

	Liquidity >0		Liquidity=0
	(1)	(2)	(3)
<i>dep var:</i>	payoff	profit	profit=payoff
Liquidity	0.740*** [0.0855]	-0.159* [0.0879]	
Session 2	183.4*** [44.03]	123.9*** [35.68]	
Session 3	283.9*** [47.04]	211.7*** [35.56]	
Session 4	250.0*** [33.52]	184.3*** [35.51]	
Session 9			58.82* [30.03]
Session 10			
Clock	65.39*** [16.36]	26.51 [24.31]	60.13*** [16.08]
_cons	-147.2** [54.83]	-93.30* [51.58]	88.66*** [16.64]
Subject FE	Yes	Yes	Yes
N	256	256	128
adj. R-sq	0.42	0.14	0.30
Robust standard errors in brackets; * p<0.1 ** p<0.05 *** p<0.01			

The second observation from Table 7 is the large influence the liquidity bonus has on payoffs. For every \$1 in liquidity bonus, payoffs are increased by a statistically significant \$0.74, while profits are reduced by an insignificant \$0.15. The positive effect on payoffs and negative effect on profits are expected as a higher liquidity bonus should motivate players to bid more aggressively on some of their securities, driving the profits negative on those securities, but securing a positive payoff with the liquidity bonus. Given that liquidity bonus is directly added to a bidder's payoff, the coefficient on the liquidity bonus can be interpreted as the percentage of the bonus captured by bidders; overall bidders captured 74% of their liquidity bonus over the four days.

Finally, Table 7 illustrates the effect of learning. Between sessions 1-3 when there was a positive liquidity bonus, we see that payoffs and profits steadily increased. Specifically, Session 2's payoffs were \$183 greater than Session 1 and Session 3's payoffs were an additional \$100 greater than Session 2. With respect to profit, Session 2's profits were \$124 greater than Session 1 and Session 3's profits were an additional \$88 greater than Session 2. In Session 4, however, the effect on learning appears to change. There is not a statistically significant difference between Session 4 and Session 3 in either the profit or payoff measure, which suggests that participants had learned all they could during the first 3 sessions. Alternatively, it could be the case that in Session 4 players were still learning, but because everyone was optimally responding to each other, there was no change in payoffs or profits.

The effect of learning is reversed when the liquidity bonus is set to zero. This is demonstrated in Regression 3. We see that players in auction pair 1 and 2 during Session 9 earned a statistically significant \$59 more than during auction pair 1 and 2 of Session 10. We consider this result below in the discussion of Table 9.

Table 8. Dependent variable = payoff; clustered standard errors

	Liquidity >0				Liquidity=0
	(1)	(2)	(3)	(4)	(5)
	payoff	payoff	payoff	payoff	payoff
Liquidity	0.740*** [0.0855]	0.810*** [0.0594]	0.810*** [0.0595]	0.875*** [0.0609]	
Session 1	omitted	omitted	omitted	omitted	
Session 2	183.4*** [44.03]	186.2*** [43.97]	186.2*** [44.06]	188.8*** [43.30]	
Session 3	283.9*** [47.04]	285.4*** [47.14]	285.4*** [47.23]	290.9*** [48.74]	
Session 4	250.0*** [33.52]	249.7*** [34.87]	249.7*** [34.94]	252.4*** [34.70]	
Session 9	0 [0]	0 [0]	0 [0]	0 [0]	65.85* [29.60]
Session 10	0 [0]	0 [0]	0 [0]	0 [0]	<i>omitted</i>
Clock	65.39*** [16.36]	65.39*** [16.39]	119.6** [31.78]	65.34*** [16.34]	59.09** [16.08]
8_Bidders		185.6*** [28.33]	239.7*** [37.68]	190.8*** [29.24]	
Clock*8_Bidders			-108.3** [35.52]		
e_payoff				-0.0914 [0.0648]	0.156 [0.116]
_cons	-147.2* [54.83]	-274.2*** [43.63]	-301.3*** [47.67]	-265.9*** [42.81]	193.9*** [30.30]
Subject FE	Yes	Yes	Yes		Yes
N	256	256	256		128
adj. R-sq	0.419	0.552	0.551		0.315
Standard errors in brackets clustered on subjects; * p<0.05 ** p<0.01 *** p<0.001					

These results are further reinforced in Table 8 where we explore the effect of competition and the expected payoff on actual payoffs. The most striking result in Table 8 is that increasing competition in both the sealed bid and clock auctions results in a higher expected payoff for all players. Using the coefficients in Table 8, the incremental payoff for the various auctions are as follows (assuming x payoff in the 4-person Sealed Bid auction):

- 8-person Sealed Bid auction: $x + \$239.70$

- 4-person Clock auction: $x + \$119.60$
- 8-person Clock auction: $x + \$251.00$

This competition benefit can be explained in the clock auction by the fact that players are learning information about all eight signals. That is, there are eight signals drawn for both auctions but in the 4-person auction four of those signals are not represented by any players. Thus it is impossible for players to learn anything about those four signals. As a result, the 8-player clock auction reveals much more information about the common value for each security than the 4-player clock auction which results in higher player payoff.

We also see that going from four to eight bidders in the sealed-bid auction increases payoffs. There is no theoretical support for this finding and thus we suggest that it is an experimental artifact. It is likely caused by the fact that four of the eight drawn signals were not observed by any player and thus all of the possible outcomes (from the eight signals) were not represented in the outcomes.

Table 8 also illustrates that there is no correlation between actual payoff and expected payoff (i.e., simulated payoff), independent of liquidity. To further explore this result, we consider the interaction of expected payoff and session-specific effects for the zero liquidity experiments in Sessions 9 and 10. These results, shown in Table 9, illustrate that while there is no correlation between actual and expected payoff during Session 9, there is a weakly-significant correlation in Session 10 (significant at the 9% threshold). Thus it appears that players deviated from the benchmark during Session 9, causing some players to experience larger profits, but reduced their deviation during Session 10, lowering average profits.

Table 9. Dependent variable = payoff; clustered standard errors

<i>Dep Var</i>	Liquidity=0	
	(1) payoff Session=9	(2) payoff Session=10
Clock	65.93* [24.91]	50.47** [15.19]
e_payoff	0.0722 [0.301]	0.297 [0.164]
_cons	146.6*** [20.87]	64.34*** [10.77]
Subject FE	Yes	Yes
N	64	64
adj. R-sq	0.385	0.633
Standard errors in brackets clustered on subjects; * p<0.05 ** p<0.01 *** p<0.001		

Tables 10 and 11 illustrate the effects of adding various additional fixed effects to the regressions presented before for the auctions with liquidity and those without liquidity, respectively. Specifically, we include competitor fixed effects, which control for the effect of playing against specific opponents in the various auctions and subject*session fixed effects, which control for subject learning over the testing period. Table 10 demonstrates that during Sessions 1-4, adding competitor fixed effects reduces the effect of liquidity in determining the payoff and dampens the effect of learning. Table 11, shows a similar phenomenon for Sessions 9-10. Adding subject*session fixed effects only appears to affect the

importance of liquidity in determining total payoff. This suggests that over time, players got better at optimally managing their liquidity bonus.

Table 10. Dependent variable = payoff for only sessions 1-4; clustered standard errors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	payoff	payoff	payoff	payoff	payoff	payoff	payoff	payoff	payoff
liquidity	0.740*** [0.0855]	0.810*** [0.0594]	0.875*** [0.0609]	0.761*** [0.0679]	0.761*** [0.0679]	0.809*** [0.110]	0.791*** [0.127]	0.891*** [0.0930]	0.928*** [0.0795]
Session 1	<i>omitted</i>	<i>omitted</i>	<i>omitted</i>						
Session 2	183.4*** [44.03]	186.2*** [43.97]	188.8*** [43.30]	385.5* [151.2]	385.5* [151.2]	397.8* [150.2]			
Session 3	283.9*** [47.04]	285.4*** [47.14]	290.9*** [48.74]	522.0** [145.6]	522.0** [145.6]	528.1** [143.8]			
Session 4	250.0*** [33.52]	249.7*** [34.87]	252.4*** [34.70]	440.7** [142.5]	440.7** [142.5]	440.8** [141.9]			
Clock	65.39*** [16.36]	65.39*** [16.39]	65.34*** [16.34]	65.39** [17.01]	65.39** [17.01]	65.35** [16.99]	65.39** [17.93]	65.39** [17.97]	65.36** [17.98]
8_Bidders		185.6*** [28.33]	190.8*** [29.24]		935.7 [914.7]	777.4 [954.3]		187.9*** [32.52]	191.1*** [32.74]
e_payoff			-0.0914 [0.0648]			-0.0737 [0.134]			-0.0574 [0.0955]
_cons	-147.2* [54.83]	-274.2*** [43.63]	-265.9*** [42.81]	-545.2*** [135.5]	156.6 [716.1]	41.62 [741.0]	-222.1** [73.70]	-372.2*** [63.99]	-291.7*** [59.08]
Subject FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Competitor FE				Yes	Yes	Yes			
Subject*Session FE							Yes	Yes	Yes
N	256	256	256	256	256	256	256	256	256
adj. R-sq	0.419	0.552	0.552	0.638	0.638	0.638	0.407	0.569	0.568

Standard errors in brackets clustered on subjects; * p<0.05 ** p<0.01 *** p<0.001

Table 11. Dependent variable = payoff for only sessions 5-6; clustered standard errors

	(1)	(2)	(3)	(4)	(5)	(6)
	payoff	payoff	payoff	payoff	payoff	payoff
Session 9	58.82 [30.03]	65.85* [29.60]	26.71 [18.06]	35.74 [20.68]		
Session 10	<i>omitted</i>	<i>omitted</i>	<i>omitted</i>	<i>omitted</i>		
Clock	60.13** [16.08]	59.09** [16.08]	60.13** [16.92]	59.34** [16.92]	60.13** [17.07]	58.65** [17.30]
e_payoff		0.156 [0.116]		0.119 [0.0799]		0.223 [0.178]
_cons	88.66*** [16.64]	75.71*** [17.51]	76.67* [34.71]	75.43* [35.26]	94.37*** [8.534]	75.72*** [15.74]
Subject FE	Yes	Yes	Yes	Yes	Yes	Yes
Competitor FE			Yes	Yes		
Subject*Session FE					Yes	Yes
N	128	128	128	128	128	128
adj. R-sq	0.299	0.315	0.649	0.657	0.447	0.471

Standard errors in brackets clustered on subjects; * p<0.05 ** p<0.01 *** p<0.001

4.2 Experiment 2: pooled securities

Table 12 reinforces the conclusions from Tables 5-7 for the Pooled security setting. The clock auction format creates no statistically significant change in profit, but does increase the payoff

significantly. That is, the clock auction format is more efficient at helping Treasury determine which banks are most in need of liquidity. Also in Table 12 we see that the bidders are able to capture slightly less (61%) of their liquidity bonus on average than in Experiment 1 (74%). Similar to Experiment 1, Table 12 demonstrates that a larger liquidity bonus caused players to bid more aggressively resulting in a statistically significant lower profit. This strategy is reflected in the individually reported strategies summarized in Appendix A.

Table 12. Experiment 2: Pooled security reverse auction

	(1)	(2)	(3)
<i>dep var:</i>	payoff	profit	payoff=profit
Liquidity	0.606*** [0.0545]	-0.119*** [0.0217]	
Session 5	<i>omitted</i>	<i>omitted</i>	
Session 6	189.9 [220.0]	-120.8 [87.43]	
Session 7	326.4 [219.8]	-225.6* [87.37]	
Session 8	-315.4 [219.9]	-791.3*** [87.42]	
Session 11			-84.66 [46.74]
Session 12			<i>omitted</i>
Clock	496.7** [150.5]	105.5 [59.82]	-84.49 [46.74]
LessPrecise	-44.68 [150.6]	-43.82 [59.87]	
Subject FE	Yes	Yes	Yes
N	256	256	128
adj. R-sq	0.364	0.388	0.265
Standard errors in brackets; * p<0.05 ** p<0.01 *** p<0.001			

The effect of learning over the first four sessions of Experiment 2 is somewhat more complicated than in Experiment 1. That is, the payoffs in Sessions 6-8 are statistically indistinguishable from Session 5. However, we observe a statistically insignificant decline in profits between Sessions 5, 6, and 7 and a dramatic and significant decline between Session 7 and Session 8 (-\$566). This decline in profits between Session 7 and Session 8 is matched by a statistically significant decline in payoffs between (decline of \$642, significant at the 99% level). This suggests that participants played significantly more aggressively during Session 8, but to their own detriment. During the days with no liquidity bonus, it appears that learning may have played a role, albeit a weak one. The payoff during Session 11 was \$85 less than Session 12 and significant at the 93% level.

Table 13 provides additional insights into the process of learning by looking at the interaction between the liquidity bonus and session and clock and session. We see an upward trend in the percentage of the liquidity bonus captured by participants over the whole of Experiment 2 when liquidity was positive. This suggests that participants became more adept at managing their liquidity

constraint over time. In addition, we see a negative trend in the benefit of clock compared to sealed-bid. This suggests that participants determined a strategy in early rounds and played that strategy independent of other bidders' actions. That is, initially when participants were unfamiliar with the pooled auction setting, the additional information revealed during the clock auction increased payoffs by a statistically significant \$1,030. This benefit of the clock declined over time independent of the presence of liquidity.

Table 13. Payoffs over time

	(1)	(2)	(3)	(4)	(5)	(6)
	Session 5	Session 6	Session 7	Session 8	Session 11	Session 12
<i>dep var:</i>	payoff	payoff	payoff	payoff	payoff	payoff
Liquidity	0.555*	0.647***	0.637***	0.727**	0	0
	[0.197]	[0.147]	[0.0769]	[0.201]	[0]	[0]
Clock	1030.1**	846.9*	224.3	-114.5	108.4**	60.57*
	[265.4]	[379.1]	[231.5]	[278.0]	[29.05]	[24.91]
LessPrecise	200.9	-266.2	-139.4	-0.490	0	0
	[407.1]	[320.5]	[228.3]	[267.6]	[0]	[0]
_cons	-39.38	-28.80	466.0	-548.4	105.8***	214.4***
	[1118.1]	[863.5]	[428.8]	[965.0]	[14.53]	[12.45]
Subject FE	Yes	Yes	Yes	Yes	Yes	Yes
N	64	64	64	64	64	64
adj. R-sq	0.625	0.564	0.749	0.655	0.443	0.61
Standard errors in brackets clustered on subjects; * p<0.05 ** p<0.01 *** p<0.001						

The final result from Table 12 is that we see no statistically significant effect with respect to the more or less precise case. That is, providing bidders with reference prices that represent 25% or 50% of the total signals does not result in a significant change in payoff. As might be expected, the more precise cases appear to result in slightly higher payoffs (\$45), but not statistically so. And when we look at the effect of the more or less precise case by day (Table 13), we again do not see a statistically significant affect or trend.

Given that there was not a tractable theoretical benchmark strategy with which to provide the auction participants, participants were forced to determine their own bidding strategies. At the conclusion of the auction all participants provided a short synopsis of their strategies (see Appendix A). Participants described strategies that were heavily determined by their liquidity draws and ratio of private signals to reference prices. Using this information, we calculated an applied bidder strategy (ABS) ratio that appears to capture the substance of how bidders used their private information. This applied bidder strategy (ABS) ratio is calculated as follows:

$$b_{is} = \frac{(sig \cdot u_{is} + (8 - sig) \cdot E[v_{is}])}{8 \cdot ref_i}$$

where *sig* is the number of private signals given to each player for security *i*, u_{is} is the average of those private signals and $E[v_{is}]$ is the expected value of the unknown signals given the known uniform distribution for securities in that pool type (75 for high quality securities and 25 for low quality securities). Finally, ref_i is the reference price for security *i* which is given to all players.

Table 14 illustrates that when the lowest ABS ratio for participants was low (i.e., ABS ratio < 0.7) bidders did significantly better than when the lowest ABS ratio was higher, independent of the presence of the liquidity bonus. For example, there is not a statistically significant difference between an ABS ratio < 0.6 and an ABS ratio between 0.6 and 0.7. However, when the ABS ratio is between 0.7 and 0.8, payoffs fall by a statistically significant \$906. When the ABS ratio rises above 0.8, payoffs are \$674 to \$780 lower than when the ABS ratio is less than 0.6.

Table 14. Effect of the lowest applied bidder strategy ratio on outcomes*

	(1)	(2)	(3)
<i>dep var:</i>	payoff	profit	payoff=profit
Liquidity	0.604*** [0.0753]	-0.121*** [0.0224]	0 [0]
Session 5	<i>omitted</i>	<i>omitted</i>	0 [0]
Session 6	150.7 [401.8]	-76.16 [114.3]	0 [0]
Session 7	357.7 [316.6]	-172.3 [130.4]	0 [0]
Session 8	-336.7 [276.8]	-729.9*** [145.4]	0 [0]
Session 11	0 [0]	0 [0]	0.893 [53.35]
Session 12	0 [0]	0 [0]	<i>omitted</i>
Clock	496.7* [199.4]	-105.5* [47.92]	84.49 [49.80]
LessPrecise	-73.33 [148.4]	42.52 [69.51]	0 [0]
ABS<0.6	<i>omitted</i>	<i>omitted</i>	<i>omitted</i>
0.6 ≤ ABS < 0.7	-546.9 [377.6]	15.88 [139.3]	-135.1 [115.1]
0.7 ≤ ABS < 0.8	-905.8** [254.3]	-105.3 [96.79]	-267.7* [111.8]
0.8 ≤ ABS < 0.9	-674.4* [261.3]	-93.38 [102.3]	-329.5** [110.3]
0.9 ≤ ABS < 1	-780.8*** [189.4]	-340.0*** [63.61]	-502.3*** [120.2]
1 ≤ ABS	-749.2** [258.7]	-528.6** [167.2]	-361.7 [223.1]
_cons	856.2 [579.2]	340.5 [180.3]	446.9*** [103.5]
Subject FE	Yes	Yes	No
N	256	256	128
adj. R-sq	0.410	0.476	0.165
Standard errors in brackets clustered on subjects when subject FE used;			
* The ABS ratio was calculated for each security and the lowest was used			

5 Recombining the sealed-bid results

Although we held only four sessions of the 8-security simultaneous reverse auction, and thus a total of four low-competition and four high-competition auctions, we can evaluate a rich set of data to determine the full range of possible outcomes from the bidding strategies employed by the subjects. To do this, we use a recombinant procedure.¹⁰ The results of our recombinant analysis add strength to one of our major contentions—the sealed-bid format results in more varied outcomes than does the clock auction format. Further, we can also assert that a few anomalous bids—mistakes—can drive the outcomes of entire auctions using the sealed-bid format. In sum, the downside risk of poor price discovery and extremely low payoffs to the bidders, are higher using the sealed-bid format.

The recombinant procedure is in principle very simple. To understand the theoretical justification, we present a simple example. Imagine a basic sealed-bid auction with two bidders. Suppose there are two of these basic auctions, auction A and auction B . Bidder 1 faces bidder 2 in auction A , while bidder 3 faces bidder 4 in auction B . A single outcome results from each auction, giving us a total of two price observations, for example. Now, consider that each bidder determined their bid using only their own private information, and no special knowledge of their opponents. Provided bidding is *anonymous*, we can expect that bidder 1 would have submitted the same bid if she had faced bidder 3 or bidder 4, just as she did when she faced bidder 2 in the auction we first observed. Exploiting this concept, we can compute the outcomes of several more auctions than those we actually observed. The total set of auctions is given by the set of bids $\{A \equiv \{b_1, b_2\}, \{b_1, b_3\}, \{b_1, b_4\}, \{b_2, b_3\}, \{b_2, b_4\}, B \equiv \{b_3, b_4\}\}$. We get a total of six outcomes to analyze instead of just two.¹¹

We use this procedure to examine more closely the outcomes of the sealed-bid simultaneous reverse auctions.¹² We observe a total of 1,820 4-bidder auctions per session, rather than 4, and a total of 12,870 8-bidder auctions per session, rather than 2. This gives us the ability to see a highly precise distribution of the outcomes that could have been produced by the strategies our subjects employed.¹³

The most important finding from this analysis is that the sealed-bid procedure is vulnerable to anomalous bids. Strong evidence is provided in Figure 1, which shows a histogram of the results in the 4-bidder case. We display the difference between true common value and price, and point out the bimodality of the data in sessions one and three. In both of these cases, small mistakes by bidders led to large differences between price and the true value of the securities. This clearly hampers price discovery, and certainly affects the profits of the winning bidders. Figure 2 shows the profit from the corresponding security sales. A heavy left tail is evident in the distribution of profits, driven again by the tendency of small mistakes by a few bidders to negatively influence the price of securities. Finally, note what happens to bidder payoffs under the same circumstances. Extremely negative outcomes are

¹⁰ See Charles and Reiley (2006) for details on the recombinant estimator.

¹¹ In this case, we get six outcomes, or four-choose-two. In general, we can calculate the number of outcomes equal to n -choose- k , where n is the total number of bidders and k is the number of bidders in each auction.

¹² Note that the theoretical justification for the recombinant procedure does not hold in a clock auction, since what one bidder does depends importantly on the signals she receives from her competitors. Thus the assumption of anonymity breaks down, and we cannot justify a recombinant analysis. Likewise, we cannot analyze reference price auctions this way, since bidders in different auctions react to information that is specific to the auction, and thus combining bids from different auctions is not theoretically or statistically valid.

¹³ Note carefully that we do not claim that all of these auctions are statistically independent. We cannot calculate multivariate test statistics with this data, as the distribution of the standard error is unknown. Instead, we use the procedure to examine in some detail what is possible, had our random matching of subjects turned out differently.

evident in session one. Even in sessions two through four, after learning has taken place, there is substantial mass at and below 200, substantially less than the average payoff of 388 we report in table 2. The average payoff statistic from the sealed-bid auctions masks the possibility of very negative outcomes. The reason is simple—when a small number of bidders drive prices below their common value, due to poor strategy or low signals, they inhibit the ability of other bidders to satisfy their liquidity needs. When liquidity needs are an especially important goal of the auction, the sealed-bid format can result in an especially bad case of allocative inefficiency.

Figure 1. Difference between value and price

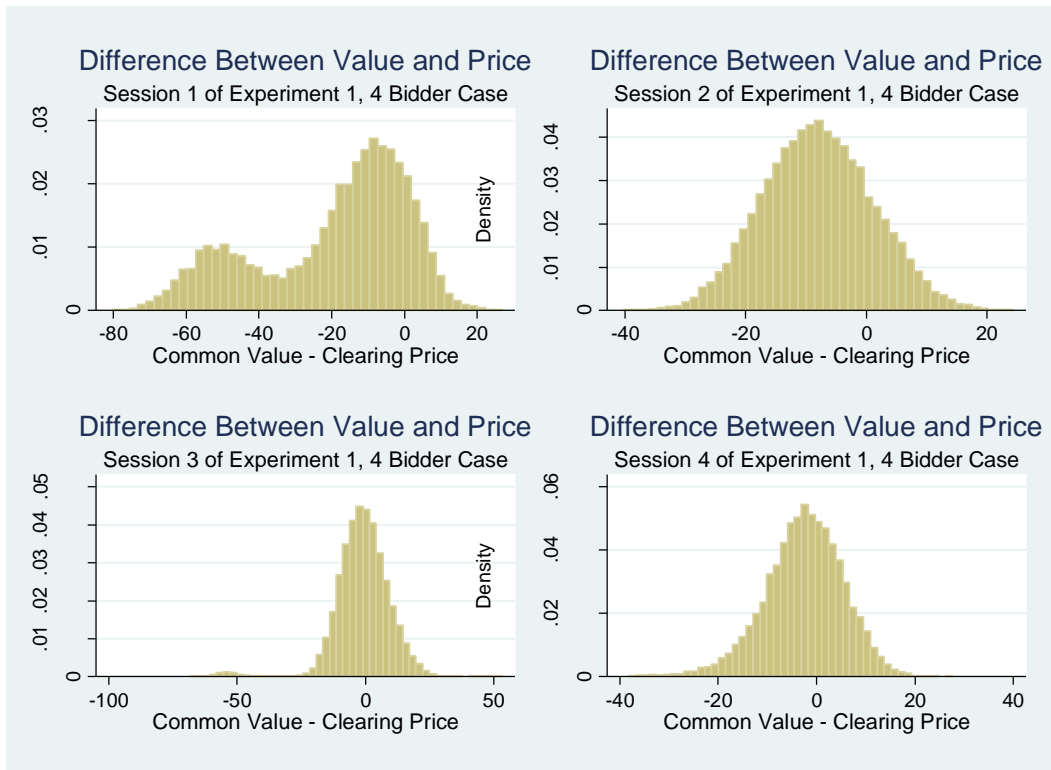

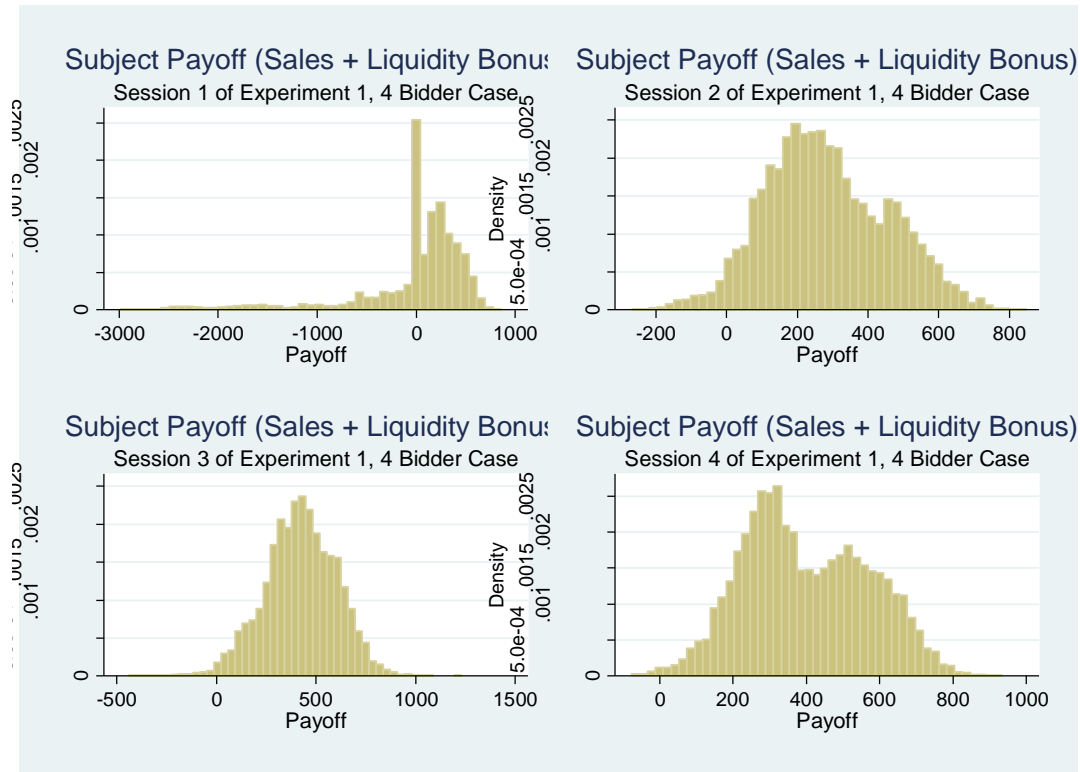


Figure 2. Profit



Such poor outcomes are extremely unlikely under the clock auction format. First, mistakes by bidders are simply less likely under the clock format. Since bidders have multiple chances to express their preferences and can continually update their strategies throughout the bidding process, anomalous bids are less likely. Second, as the results of our econometric analysis demonstrate, bidders are better able to manage their liquidity needs under the clock format. Thus, should a small number of bidders drive the price of a handful of securities well below their common value, the clock format enables other bidders to respond by adjusting their bidding on other securities. Sensible bidders will bid more aggressively on other securities in their portfolio in order to meet their liquidity needs, allowing them to recover reasonable payoffs. The sealed-bid format gives bidders no such chance to adjust their strategies. When liquidity needs are a dominant concern of the Treasury, the sealed-bid format leaves gains from trade on the table.

Figure 3. Payoffs



6 Analysis and discussion

6.1 Feasibility of implementation

One of the points in conducting the experiments was to demonstrate the feasibility and practicality of conducting a computerized auction in which multiple securities are purchased simultaneously. In this regard, both the sealed-bid and descending-clock auctions can be regarded as successfully implemented in a short time.

6.2 Competitiveness of price

Notwithstanding the presence of adverse selection, as a theoretical matter, the price in the auction can be driven below the security's fundamental value to the extent that the bidders have liquidity needs above and beyond their objective of earning trading profits. This was demonstrated in the experiments. In both the sealed-bid and descending-clock auctions, the prices were significantly below the fundamental values of the respective securities. This is seen in the second row (profit) in Tables 3-6. A bidder's mean trading profits in Experiment 1 were negative, though not significantly so under the sealed-bid (-39) and clock (-13). The mean trading profits in Experiment 2 were significantly negative under both formats: -693 with the sealed-bid format and -799 with the clock format.

6.3 Frequency of erroneous bids

In a relatively small but clearly noticeable subset of the auctions, bidders submitted what were almost certainly unintentionally low bids ("erroneous bids"). In descending-clock auctions, the harm to a bidder from an erroneous bid was often fairly limited, as the extent of the damage was that the bidder found herself still in the auction at the end of the round. However, in sealed-bid auctions, the harm

could be much greater. Thus, one advantage observed of the clock auction was that it helped to insulate bidders from the effects of their own mistakes.

6.4 Management of liquidity needs

If separate sealed-bid auctions are conducted simultaneously, the bidder has no way to condition her bidding in one auction on the amount that she is likely to win in other auctions. By contrast, if separate dynamic auctions are conducted simultaneously, the bidder can observe the outcome evolving in one auction and use that information to adjust her behavior in other auctions. Thus, the bidder has much greater ability to manage her liquidity need, without “overshooting” and selling more securities at prices below value than the bidder needs to sell.

This is demonstrated clearly in the experimental results. Despite that the bidders had the same values and the same liquidity needs in the sealed-bid auctions and in the clock auctions, the bidders attained average payoffs in Experiment 1 (Experiment 2) of 388 (3,222) in the sealed-bid auctions and of 453 (3,719) in the clock auctions, as shown in row 5 of Tables 3-6. The payoffs in the clock auctions were significantly higher (at the 1% level) than in the sealed bid auction. This is entirely due to better management of liquidity with the clock auction. As shown in the third row of Tables 3-6, the mean liquidity bonus in Experiment 1 (Experiment 2) was significantly larger, 466 (4,517) with the clock format, compared with 428 (3,915) with the sealed bid format. Under sealed-bid, the bidder often overshoots her liquidity need. The mean overshoot in Experiment 1 (Experiment 2) with sealed-bid was 692 (1,984), compared with 605 (905) for the clock format (see row 7 of Tables 3-6). This difference is significant at the 1% level (p-value of 0.0014).

6.5 Cost of purchasing securities

There appear to be three distinct effects determining the comparison of purchase prices between the sealed-bid and dynamic auction formats. First, a dynamic auction format is generally known for mitigating the winner’s curse, leading to more aggressive bidding and to lower prices in a reverse auction. Second, as seen above, the bidders submit fewer erroneous bids in a dynamic format, leading to higher prices. Third, as seen above, a dynamic format allows bidders to manage better their liquidity needs; more than likely, this would lead to fewer desperation offers and cause higher prices. Combining these three effects, the price comparison between sealed-bid and dynamic auction formats is ambiguous.

In the experimental results, the price difference between the two formats is not statistically significant. The bidders, with the same values and the same liquidity needs in the two auction formats produced mean clearing prices of 49.57 in the clock and 47.79 in the sealed-bid auction in Experiment 1. In Experiment 2, the two formats produced mean clearing price points of 85.22% in the sealed-bid auctions and of 83.87% in the clock auctions (see row 1 of Tables 3-6).

Combining the results on satisfying liquidity needs and on the cost of purchasing securities, the taxpayer would favor the descending-clock auction. While the taxpayer’s expenditure is approximately equal in the sealed-bid and the clock auction, the clock auction gives “more bang for the buck”—for a given expenditure of money, the clock auction better directs resources towards satisfying the liquidity needs of the banking sector.

6.6 Variability of outcomes and informativeness of prices

Finally, there appears to be a fundamental difference in the experimental results between the sealed-bid auction and the corresponding dynamic auction. All other things being equal, the outcomes of the sealed-bid auction are more variable and random.

First, this means that the outcomes of the dynamic auction are more predictable, and thus more satisfying to risk-averse banks and regulators. The greater variability is seen in the variance of the bidder's payoff. The standard deviation of the bidder's payoff in Experiment 1 (Experiment 2) is 222 (1,312) with the clock format compared with 285 (1,654) with sealed-bid (see row 6 of Tables 3-6). This difference is significant at the 1% level (p -value of 0.0095).

Second, one of the objectives of government purchases of troubled securities is to restart frozen secondary markets by providing relevant transaction prices. By doing so, the government can rely on the private market to accomplish some of the government's objectives, reducing the need for government resources. The experimental design limited the extent to which this can be seen in the data, as there were only two separate pools of securities and thus only two independent prices determined in the auctions. Nevertheless, it can be seen that the prices resulting from the dynamic auction are more accurate, an effect which can be expected to be enlarged when more independent prices are determined by an auction.

7 Conclusion

We present our findings from laboratory experiments of troubled asset reverse auctions. The experiments demonstrate the feasibility of implementing a purchase program for troubled assets in a short period of time using either a sealed-bid or a dynamic auction format. The experiments suggest that the taxpayer cost of purchases using a well-designed and well-implemented auction program could be small using either auction format, to the extent that sellers have substantial liquidity needs. However, the dynamic auction format has significant advantages over the sealed-bid auction format for both the banks and the taxpayers, because the informational feedback provided during the auction enables the seller to manage better its liquidity needs.

Our experiments focused on trading profits and liquidity needs as the bidder's principal objectives. In practice, a seller of troubled assets also cares about its portfolio risk. For reasons of simplicity, we ignored portfolio risk in the experiments. However, there is good reason to expect that what we learned about a bidder's challenges in managing liquidity would carry over to the issue of portfolio risk. It is the bidder's ability to see, while the auction is still running, which asset sales are likely to be successful that enables the bidder to better reach its revenue target in the dynamic auction. By the same token, the bidder's ability to see which asset sales are likely to be successful should enable the bidder to better manage its portfolio risk. Thus, explicitly including portfolio risk would likely strengthen the case for dynamic auctions over sealed-bid auctions for purchases of troubled assets.

References

- Alsemgeest, Paul, Charles Noussair, and Mark Olson (1996), "Experimental Comparisons of Auctions Under Single- and Multi-Unit Demand," Working Paper, Purdue University.
- Ausubel, Lawrence M. (2004), "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94:5, 1452-1475.
- Ausubel, Lawrence M. (2006), "An Efficient Dynamic Auction for Heterogeneous Commodities," *American Economic Review*, 96:3, 602-629.
- Ausubel, Lawrence M. and Peter Cramton (2002), "Demand Reduction and Inefficiency in Multi-Unit Auctions," University of Maryland Working Paper 96-07, revised July 2002.
- Ausubel, Lawrence M. and Peter Cramton (2004), "Auctioning Many Divisible Goods," *Journal of the European Economic Association*, 2, 480-493, April-May.

- Ausubel, Lawrence M. and Peter Cramton (2006), "Dynamic Auctions in Procurement," Dimitri, Nicola, Gustavo Piga, and Giancarlo Spagnolo, *Handbook of Procurement*, Cambridge University Press, 220-245.
- Ausubel, Lawrence M. and Peter Cramton (2008a), "A Troubled Asset Reverse Auction," Working Paper, University of Maryland.
- Ausubel, Lawrence M. and Peter Cramton (2008b), "Auction Design Critical for Rescue Plan," *Economists' Voice*, 5:5, www.bepress.com/ev/vol5/iss5/art5.
- Ausubel, Lawrence M. and Peter Cramton (2009), "A Two-Sided Auction for Legacy Loans," University of Maryland.
- Barnett-Hart, Anna Katherine (2009), "The Story of the CDO Market Meltdown: An Empirical Analysis," Harvard College Department of Economics.
- Brunnermeier, Markus K. (2009). "Deciphering the Liquidity and Credit Crunch 2007-2008," *Journal of Economic Perspectives*, 23(1): 77-100.
- Congressional Budget Office. (2008). "Policy Options for the Housing and Financial Markets."
- Cramton, Peter (1998), "Ascending Auctions," *European Economic Review*, 42, 745-756.
- Cramton, Peter, Yoav Shoham, and Richard Steinberg (2006), *Combinatorial Auctions*, Cambridge, MA: MIT Press.
- French, Kenneth R., Martin N. Baily, John Y. Campbell, John H. Cochrane, Douglas W. Diamond, Darrell Duffie, Anil K. Kashyap, Frederic S. Mishkin, Raghuram G. Rajan, Robert J. Shiller, Hyun Song Shin, Matthew J. Slaughter, Jeremy C. Stein, and René M. Stulz (2010). *The Squam Lake Report: Fixing the Financial System*. Princeton, NJ: Princeton University Press.
- Gorton, Gary B. (2008). "The Panic of 2007," Yale ICF Working Paper No. 08-24.
- Gorton, Gary (2009). "Information, Liquidity, and the (Ongoing) Panic of 2007," *American Economic Review*, 99(2): 567-72.
- Kagel, John H. (1995), "Cross-Game Learning: Experimental Evidence from First-Price and English Common Value Auctions," *Economic Letters*, 49, 163-170.
- Kagel, John H., Ronald M. Hartstad, and Dan Levin (1987), "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55, 1275-1304.
- Kagel, John H. and Dan Levin (2001), "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions," *Econometrica*, 69, 413-454.
- Kagel, John H. and Dan Levin (2005), "Multi-Unit Demand Auctions with Synergies: Some Experimental Results," *Games and Economic Behavior*, 53, 170-207.
- Kagel, John H. and Dan Levin (2008), "Implementing Efficient Multi-object Auction Institutions: An Experimental Study of the Performance of Boundedly Rational Agents," *Games and Economic Behavior*, forthcoming.
- Klemperer, Paul (2009). "A New Auction for Substitutes: Central-Bank Liquidity Auctions, 'Toxic Asset' Auctions, and Variable Product-Mix Auctions," CEPR Discussion Paper No. DP7395.
- Levin, Dan, John H. Kagel, and Jean-Francois Richard (1996), "Revenue Effects and Information Processing in English Common Value Auctions," *American Economic Review*, 86, 442-460.
- McAdams, David (2007), "Adjustable Supply in Uniform Price Auctions: Non-Commitment as a Strategic Tool," *Economics Letters*, 95, 48-53.
- McAfee, R. Preston and John McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25, 699-738.
- McCabe, Kevin A., Stephen J. Rassenti, and Vernon L. Smith (1990), "Auction Institutional Design: Theory and Behavior of Simultaneous Multiple-Unit Generalizations of the Dutch and English Auctions," *American Economic Review*, 80, 1276-1283.

- Mizen, Paul (2008), "The Credit Crunch of 2007-2008: A Discussion of the Background, Market Reactions, and Policy Responses," *Federal Reserve Bank of St. Louis Review*, 90(5): 531-568.
- Milgrom, Paul (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
- Milgrom, Paul and Robert J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089-1122.
- Mullin, Charles H and David H Reiley (2006). "Recombinant Estimation for Normal-Form Games, with Applications to Auctions and Bargaining," *Games and Economic Behavior*, 54: 159 - 182.
- Swagel, Phillip (2009). "The Financial Crisis: An Inside View," *Brookings Papers on Economic Activity*, Spring 2009, p. 1-63.

Appendix A

Analysis of Bidding Strategy in Experiment 2

Each of the 19 expert bidders provided a summary of their bidding strategy during Experiment 2. Although there are many variations of strategies employed, the following summarize the primary strategy components from the bidders.

Liquidity. All of the bidders described their strategy as heavily dependent on their liquidity draw. That is, they recognized that even though the securities were likely to close at a price below their respective common values, they could make money on the liquidity bonus. As a result, they selected certain securities on which to bid aggressively. Once determined, bidders would instantly supply-reduce to a level that just allowed them to meet their liquidity bonus if the price were to fall to their estimated dropout ratio. The immediate reduction was an effort to signal to other bidders that they should be reducing supply also. They will hold that level of the security until the price goes below their estimated dropout ratio, at which point they'd rapidly drop out of the auction. Once the liquidity bonus is satisfied they will drop out of the auction completely.

Reference Prices. Bidders were generally aware that securities with low reference prices were least likely to sell and thus they reduced their supply on those first. DI stated "I was more aggressive on the securities with a higher than average ... reference price." Many bidders compared their own signals to these reference prices as an indication of which securities to supply reduce first. Once determined, these ratios were often ranked from lowest (best) to highest (worst) to determine the appropriate aggressiveness to be used for each security. The lower the ratio on a security, the more aggressively the holder bid (i.e., the longer they held on to the security). If the ratio was above the reference price, they dropped out early.

High Quality / Low Quality. Five bidders stated that they were going to focus their attention on the high quality securities as that would be the easiest way to achieve their liquidity bonus. However, there were 3 bidders who stated that they spread their attention across both pools and one bidder who focused on the lower quality pool because he thought it "will have relatively higher or positive P-CV."

Learning. Four of the bidders stated that it took 1-2 days to optimize their strategy. The bidders were divided on whether observing the actions of the other bidders was helpful in optimizing their own behavior. Some stated that the actions of the other bidders were not a helpful signal in determining their values, either because the other bidders were making mistakes or misrepresenting. AP states that he couldn't get information from other bidders because "it was too easy to intentionally send out "bad" information...make mistakes ... and/or play irrationally." Others felt the opposite. BW said and DI concurred that they'd watch what other bidders were selling and sell the same "because that specific security may have a lower common value." Still other bidders stated that while specific information was hard to determine from the actions of other bidders, they were able to learn the general behavior of the other bidders as the auctions progressed (i.e., that bidders were more aggressive than they thought initially). KD stated, "I learned that broadly speaking, there were two types of players: those who appreciated the benefits of supply reduction and those who did not." KD continued that she would learn in the first few rounds the type of opponents she was up against and would adjust her strategy accordingly. PS stated, "I believed other bidders will [supply reduce]. I was wrong...So in the remaining session, I just hung in the securities ...without supply reducing."

Estimated Dropout Ratio. To achieve their respective liquidity goals, bidders would hold the necessary quantity of securities up to their calculated dropout ratio. These dropout ratios were estimated in two steps. First, they would calculate a target CV-ref price ratio and then they would adjust it to determine the ratio at which they would drop out. Four of the bidders calculated the target ratio assuming the expected values for the unknown signals (75 and 25 for high and low, respectively). Seven bidders used only the ratio of their private signals to the reference price. Three bidders used the clearing ratio from previous similar experiments. Once these ratios were determined, seven of the bidders stated that they'd drop out completely when the price got to 50% of the target ratio. Other bidders used different heuristics or downward adjustments as to when they'd drop out. For example, AV kept significant quantities of securities "until we got around 96-94% of the initial price." Several bidder stated that their level of confidence in their final dropout ratio was determined by the number of private signals they had: 4 private signals made them more confident with respect to the ratio; 1 signal made them less confident. Lower confidence with respect to the ratio caused them to behave as if the ratio were slightly above their calculated value.

Supply Reduction. Regardless of the reason, most bidders reduced their supply as quickly as possible, especially for those securities where their estimated true ratio was high, assuming these would not clear. Even in cases when they thought they could win, bidders often stated that they'd immediately supply reduce 50% and then hold on that quantity for a while. Many of the bidders stated that they were often hoping for tacit collusion, even when it did not come. For example, AV complains that "the other players would not reduce quantities significantly from a round to another." JR states, "What was more frustrating was the fact that people just wouldn't drop [their supply]." JB says, "I didn't see any way to elicit cooperation to get others to drop out sooner." KM says, "my aim was always to induce people to drop quantities early on by dropping myself...usually I assumed that players would take equal responsibility for reductions." Some bidders stated that they had a preference to supply reduce less for their larger holdings, because they knew more about those holdings.

Sealed Bid vs. Clock Auctions. Many of the bidders stated that it was more difficult to play the sealed bid because they weren't learning anything about the values of other bidders. As a result, they had to guess an optimal dropout ratio. BW and DI guessed an 80% ratio to drop out. Several players stated that they played more conservatively in the Sealed Bid auction.

And there were other things that were conspicuously absent from their strategies:

Differentiation between Less/More Precise. Almost no one stated a strategy difference between the two auction types. In fact, two of the bidders explicitly stated that they did nothing different between the more precise and less precise scenarios. Anyone using the bidding tool calculated the estimated true ratio with different assumptions for each of these scenarios (i.e., 12 vs. 8 total bidders for the more and less scenarios, respectively), but many bidders stated that they did not use the bidding tool.

Activity Points. Only one bidder discussed a strategy that involved shifting supply to high signal / low reference price securities within the same pool in an effort to distract other bidders. Other than this, no one talked about using the activity point constraint as a means of passing information.

Appendix B

A Common-Value Auction with Liquidity Needs: Bidder Instructions for Experiment 1

12 October 2008

In this experiment, you are a bidder in a series of auctions conducted over four sessions. Each session is held on a different day and consists of four different auctions. Bidders are randomly assigned to each auction. Thus, you do not know who the other bidders are in any of your auctions. You will be bidding from your private cubical, which includes a computer and a bidder tool that is unique to you. We ask that you refrain from talking during the experiment. If you need assistance at any time, just raise your hand and one of the experimental staff will assist you.

You will be paid at the end of the experiment. Your payment is proportional to your total experimental payoff—the sum of your payoffs from each of your auctions. In particular, you will receive *\$1 in take-home pay for every 10 million experimental dollars that you earn*. Throughout, dollar figures refer to your experimental payoff unless explicitly stated otherwise—and the “millions” are generally suppressed. When explicitly stated, your real dollar payment will be referred to as your *take-home payment*. We anticipate that each of you will earn a take-home payment of about \$100 per experimental session on average. However, your actual take-home payment will depend on your bidding strategy, the bidding strategies of the other bidders you face, and the particular realizations of numerous random events.

In each auction, you will compete with other bidders to sell your holdings of eight securities to the Treasury. The eight different securities have different values. However, for each security, the bidders have a common value, which is unknown. Each bidder has an estimate of the common value. You profit by selling securities to the Treasury at prices above the securities’ true values. You also have a need for liquidity (cash). The sale of securities to the Treasury is your source of liquidity. Thus, you care about both profits and liquidity (your sales to the Treasury).

Two formats are used:

- *Simultaneous uniform-price sealed-bid auction (“sealed-bid auction”)*. Bidders simultaneously submit supply curves for each of the eight securities. Supply curves are non-decreasing (i.e. upward-sloping) step functions. The auctioneer then forms the aggregate supply curve and crosses it with the Treasury’s demand. The clearing price is the lowest-rejected offer. All quantity offered below the clearing price is sold at the clearing price. Quantity offered at the clearing price is rationed to balance supply and demand, using the proportionate rationing rule.
- *Simultaneous descending clock auction (“clock auction”)*. The securities are auctioned simultaneously. There is a price “clock” for each security indicating its tentative price per unit of quantity. Bidders express the quantities they wish to supply at the current prices. The price is decremented for each security that has excess supply, and bidders again express the quantities they wish to supply at the new prices. This process repeats until supply is made equal to demand. The tentative prices and assignments then become final. Details of the design are presented in Ausubel and Cramton (2008), which you received earlier.

In each session, you will participate in four different auctions in the following order:

1. 4-bidder sealed-bid. A sealed bid auction with four bidders.

2. 8-bidder sealed-bid. A sealed bid auction with eight bidders.
3. 4-bidder clock. A clock auction with four bidders.
4. 8-bidder clock. A clock auction with eight bidders.

Each session, one of your two 4-bidder auctions and one of your two 8-bidder auctions will be selected at random. Your take-home payment from the session will be based on the number of (million) laboratory dollars that you earn in these two auctions only.

1 Securities

In each auction, eight securities are purchased by the Treasury. The bidders are symmetric, before the draws of bidder-specific private information about security values and liquidity needs.

In each session, two sets of bidder-specific private information are drawn independently from the same probability distributions. The first set is used in the 4-bidder auctions (auctions 1 and 3); the second set is used in the 8-bidder auctions (auctions 2 and 4). You are given no feedback following the sealed-bid auctions; thus, your behavior in the clock auctions cannot be influenced by outcomes in the sealed-bid auctions of a session.

In each 4-bidder auction, the Treasury demand is 1000 shares of each security, where each share corresponds to \$1 million of face value. Each bidder has holdings of 1000 shares of each security. Thus, it is possible for a single bidder to fully satisfy the Treasury demand for a particular security; that is, for each security there may be just a single winner or there may be multiple winners. One quarter of the total available shares will be purchased by the Treasury.

In each 8-bidder auction, the Treasury demand is 2000 shares of each security, where each share corresponds to \$1 million of face value. Each bidder has holdings of 500 shares. Thus, at least four bidders are required to fully satisfy the Treasury demand—there must be at least four winners. One half of the total available shares will be purchased by the Treasury.

2 Your preferences

From each auction, your payoff depends on your profits and how well your liquidity needs are met.

Common Value Auction

The value of each security in cents on the dollar is the average of eight iid random variables uniformly distributed between 0 and 100:

$$v_s = \frac{1}{8} \sum_{i=1}^8 u_{is}, \text{ where } u_{is} \stackrel{iid}{\sim} U[0,100].$$

Suppose you are bidder i .

Your private information about security s is the realization u_{is} . Notice that both for the 8-bidder and 4-bidder auctions, the common value is the average of eight uniform draws, so that only the first four draws are revealed in the 4-bidder auction. This means that the true values have the same distribution in both 4-bidder and 8-bidder auctions and your private information has the same precision.

If you sell the quantity q_s of the security s at the price p_s , then your profit (in million \$) is

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_{s=1}^8 (p_s - v_s) q_{is}.$$

Liquidity need

You have a liquidity need, L_i , which is drawn iid from the uniform distribution on the interval [250, 750]. You know your own liquidity need, but not that of the other bidders. You get a bonus of \$1 for every dollar of sales to the Treasury up to your liquidity need of L_i . You do *not* get any bonus for sales to the Treasury above L_i . Thus, your bonus is

$$\min \left[L_i, \frac{1}{100} \sum_{s=1}^8 p_s q_{is} \right].$$

Your payoff from an auction

Combining your profit and your liquidity bonus results in your total payoff

$$U_i(p, q_i, v) = \begin{cases} \frac{1}{100} \sum_{s=1}^8 (2p_s - v_s) q_{is} & \text{if } \frac{1}{100} \sum_{s=1}^8 p_s q_{is} < L_i \\ L_i + \frac{1}{100} \sum_{s=1}^8 (p_s - v_s) q_{is} & \text{otherwise} \end{cases}$$

Thus, an additional dollar of cash is worth two dollars when your liquidity need is not satisfied, but is worth one dollar when your liquidity need is satisfied.

3 Bidder tool and auction system

You have an Excel workbook that contains your private information for each auction. You will use the tool to submit bids in the sealed-bid auctions. In addition, the tool has features that will help your decision making in each of the auctions. Each auction has its own sheet in the tool. It is essential that you are working from the correct sheet for each auction. For example the 4-bidder sealed-bid auction is the sheet named *Sealed Bid 4*.

For the clock auctions, bidding is done via a commercial auction system customized to this setting. You use the web browser to connect to the auction system. For each clock auction, you must go to a new auction site. The URL for the auction site is given in the bidder tool on the particular auction sheet, Clock 4 or Clock 8, for the 4-bidder or 8-bidder clock auction. Once at the correct auction site, log in with the user name and password given on your auction sheet.

The auction system is easy to use. You will have an opportunity to use it in the training seminar.

An important feature of the tool is the calculation of expected security values conditional on information you specify. In a common value auction, it is important for you to condition your estimate of value on your signal *and the information winning conveys*. Since your bid is only relevant in the event that you win, you should set your bid to maximize your payoff in that event. In this way, you avoid the winner's curse, which in this case is the tendency of a naïve bidder to lose money by selling shares at prices below what they are worth. In addition, in the clock auctions, the bidder also must condition on any information revealed through the bidding process. The tool provides one flexible method of calculating an appropriate conditional expected value for each security.

4 Bidding strategy

The auction environment has three complicating features:

- *Common value auction.* You have an imperfect estimate of the good's common value.
- *Divisible good auction.* Your bid is a supply curve, specifying the quantity you wish to sell at various prices.
- *Liquidity need.* You have a specific liquidity need that is met through selling shares from your portfolio of eight securities.

The combination of these factors makes a complete equilibrium analysis impossible. Nonetheless, equilibrium analysis is possible in a simplified environment. To aid your thinking about strategy, we discuss a particular strategy, which abstracts from the complications of a divisible good auction and the liquidity needs. In particular, we assume:

1. Each bidder submits a flat supply schedule; that is, the bidder offers to sell all of her holdings of a particular security at a specified price.
2. Each bidder ignores her liquidity need, bidding as if $L_i = 0$.

With these assumptions it is possible to calculate an equilibrium strategy, which we call the *benchmark strategy*. The analysis of this strategy will be helpful in thinking about the common value feature of the auction environment.

We wish to emphasize that the benchmark strategy focuses on only one element of the auction. Your challenge is to determine your own strategy to maximize your payoff that reflects all aspects of the auction environment.

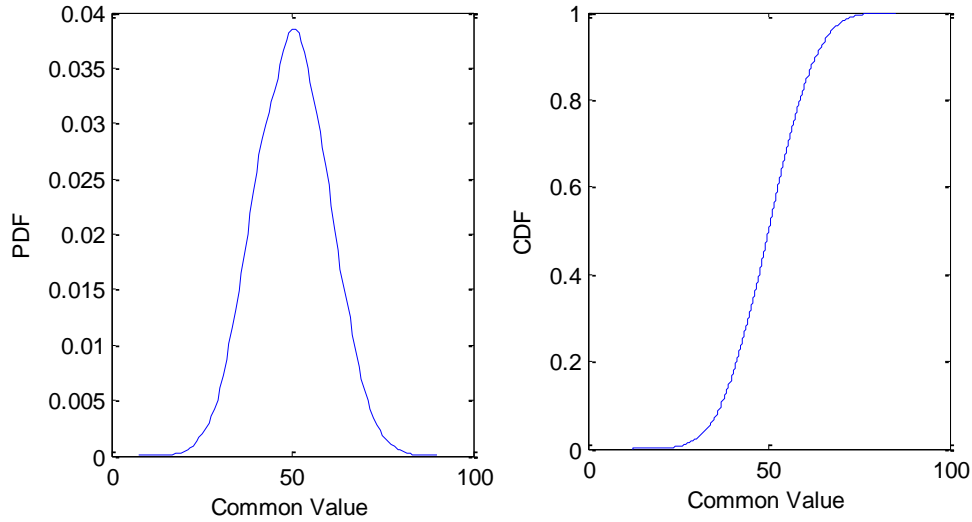
4.1 Common value distribution

As mentioned earlier, the value of each security in cents on the dollar is the average of eight iid random variables uniformly distributed between 0 and 100:

$$v_s = \frac{1}{8} \sum_{i=1}^8 u_{is}, \text{ where } u_{is} \stackrel{iid}{\square} U[0,100].$$

The pdf and cdf of the common value are shown in Figure 1. The common value has a mean of 50 and a standard deviation of 10.2. Notice that the common value is approximately normally distributed, since it is the sum of many independent draws.

Figure 1. Probability density and cumulative distribution of common value



You are bidder i . Your private information about security s is the realization u_{is} .

4.2 Sealed-bid uniform-price reverse auction

Under our strong simplifying assumptions, the auction is equivalent to a unit-supply common value auction with uniform pricing. In this case, just as in a single-item second-price auction, your bid does not determine what you pay, only the likelihood of winning, thus the best strategy is to bid your true conditional expected value for the good. The trick, however, is to condition on your signal and the information that winning conveys.

In the 4-bidder auction, under the benchmark assumptions, there is only a single winner, so the correct condition is as derived in Milgrom and Weber (1982). Your bid is your expected value conditional on your signal being the lowest and on the second-lowest signal being equal to yours:

$$b_{is} = E[v_s | u_1 = u_{is}, u_2 = u_1],$$

where u_1 is the lowest signal and u_2 is the second-lowest signal. The reason you condition on the second-lowest signal being as low as yours is that the bid must be optimal when it is on the margin and thus impacts whether you win. Your bid becomes marginal and hence decisive only in the event that you tie with the second-lowest.

For the 8-bidder auction, there are exactly four winners. Hence, we need to generalize the above formula to multiple winners. The optimal rule is to bid the expected value conditional on your signal being the fourth-lowest signal and on the fifth-lowest signal being equal to yours:

$$b_{is} = E[v_s | u_4 = u_{is}, u_5 = u_4],$$

where u_4 is the fourth-lowest signal and u_5 is the fifth-lowest signal.

Since the signals are uniformly distributed, it is easy to calculate the above conditional values. In both case, the conditional value is a linear function of your signal.

In the 4-bidder auction, there is a single winner. The conditioning is that you win, the next lowest bidder has your value, and the remaining two bidders are evenly distributed above your value:

$$\text{4-bidder sealed-bid strategy: } b_{is} = \frac{1}{8} \left(2u_{is} + 2 \left(\frac{u_{is} + 100}{2} \right) + 4 \cdot 50 \right) = \frac{3}{8} u_{is} + \frac{75}{2}.$$

In the 8-bidder auction, there are four winners. The conditioning is that you have the fourth-lowest signal, the fifth-lowest signal is the same, the three signals below you are evenly distributed below your signal and zero, and the remaining three bidders are evenly distributed above your signal:

$$\text{8-bidder sealed-bid strategy: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + 3 \left(\frac{u_{is} + 100}{2} \right) \right) = \frac{5}{8} u_{is} + \frac{75}{4}.$$

4.3 Descending clock auction

In the clock auction, under the benchmark assumption, you will observe the price at which other bidders drop out. This provides additional information on which to condition your bid. Here we assume that the price clock is continuous. In the actual experiment, the price clock is discrete, and although bidders can make reductions at any price, you will only learn the aggregate supply at the end of round price. You may want to assume the quantity reduction occurred half-way between the prior price and the ending price.

When the clock starts, you calculate your dropout point in the same way as above. As the price clock falls, one of the other bidders may drop out. When the first bidder drops out, you can calculate this bidder's draw from the following equation.

$$u_{8s} = \frac{P_{\text{dropout}} - \text{intercept}}{\text{slope}}$$

where the slope and intercept are taken the formulas above. With this new information on which to condition your bid, the revised optimal bid for the 8-bidder clock auction is straightforward to calculate.

8-bidder clock strategy

$$\text{No one has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + 3 \left(\frac{u_{is} + 100}{2} \right) \right).$$

$$\text{One bidder has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + u_8 + 2 \left(\frac{u_{is} + u_8}{2} \right) \right).$$

$$\text{Two bidders have dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + u_8 + u_7 + 1 \left(\frac{u_{is} + u_7}{2} \right) \right).$$

$$\text{Three bidders have dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 3 \frac{u_{is}}{2} + u_8 + u_7 + u_6 \right).$$

Similarly, we can calculate the analogous formulas for the 4-bidder clock auction.

4-bidder clock strategy

$$\text{No one has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 2 \left(\frac{u_{is} + 100}{2} \right) + 4 \cdot 50 \right)$$

$$\text{One bidder has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + u_4 + 1 \left(\frac{u_{is} + u_4}{2} \right) + 4 \cdot 50 \right).$$

$$\text{Two bidders have dropped out: } b_{is} = \frac{1}{8} (2u_{is} + u_4 + u_3 + 4 \cdot 50).$$

Note that the above formulas are all linear in the dropout price, so it is easy to invert to get compute the bidder's signal.

4.4 Moving beyond the benchmark strategy

The bidding tool is set up to calculate the conditional expected values assuming the benchmark strategy. Of course, you (and others) may well deviate from the benchmark strategy as a result of liquidity needs or other reasons, since these other factors are ignored in the benchmark calculation.

The bidding tool allows for this variation in a number of ways. First, in the sealed-bid auctions your bid can be any upward sloping step function to account for liquidity and possible supply reduction by you or others. Second, in the clock auctions, the tool lets you interpret what a dropout is and where it occurs. This is useful and necessary when bidders make partial reductions of supply. In addition, although the tool will calculate a particular signal based on a dropout, you are free to type in any signal inference you like. Whatever you type as "my best guess" will be used in the calculation of the conditional expected value.

Further details of the tool will be explained in the training seminar.

If you have any questions during the experiment, please raise your hand and an assistant will help you.

Remember your overall goal is to maximize your experimental payoff in each auction. You should think carefully about what strategy is best apt to achieve this goal.

Many thanks for your participation.

Appendix C

A Common-Value Reference-Price Auction with Liquidity Needs: Bidder Instructions for Experiment 2

20 October 2008

In this experiment, you are a bidder in a series of auctions conducted over four sessions. Each session is held on a different day and consists of four different auctions. Bidders are randomly assigned to each auction. Thus, you do not know who the other bidders are in any of your auctions. You will be bidding from your private cubical, which includes a computer and a bidder tool that is unique to you. We ask that you refrain from talking during the experiment. If you need assistance at any time, just raise your hand and one of the experimental staff will assist you.

You will be paid at the end of the experiment. Your payment is proportional to your total experimental payoff—the sum of your payoffs from each of your auctions. In particular, you will receive *\$1 in take-home pay for every one hundred thousand experimental dollars (100 \$k) that you earn*. Throughout, dollar figures refer to your experimental payoff unless explicitly stated otherwise—and the “thousands” are generally suppressed. When explicitly stated, your real dollar payment will be referred to as your *take-home payment*. We anticipate that each of you will earn a take-home payment of about \$100 per experimental session on average. However, your actual take-home payment will depend on your bidding strategy, the bidding strategies of the other bidders you face, and the particular realizations of numerous random events.

In each auction, you will compete with other bidders to sell your holdings of eight securities to the Treasury. The eight securities are split into two pools: four securities are low quality and four are high quality. The eight different securities have different values. However, for each security, the bidders have a common value, which is unknown. Each bidder has an estimate of the common value. You profit by selling securities to the Treasury at prices above the securities’ true values. You also have a need for liquidity (cash). The sale of securities to the Treasury is your source of liquidity. Thus, you care about both profits and liquidity (your sales to the Treasury). The Treasury has allocated a particular budget for the purchase of each pool of securities within each auction. Its demand for high-quality securities is distinct from its demand for low-quality securities.

Before each auction, the auctioneer assigns each security a *reference price* (expressed in cents on the dollar of face value), which represents the Treasury’s best estimate of what each security is worth. For example, a high-quality security might be assessed to be worth 75 cents on the dollar, while a low-quality security might be assessed to be worth 25 cents on the dollar. A *reference-price auction* determines the price-point—a percentage of the reference price—for each pool of securities. A winning bidder is paid the pool’s price-point \times the security’s reference price for each unit of the security sold.

Two formats are used:

- *Simultaneous uniform-price sealed-bid auction (“sealed-bid auction”)*. Bidders simultaneously submit supply curves for each of the eight securities. Supply curves are non-decreasing (i.e. upward-sloping) step functions, offering a quantity at each price-point. The auctioneer then forms the aggregate supply curve and crosses it with the Treasury’s demand. This is done separately for each pool (i.e., for high- and low-quality securities, separately). The clearing price-point is the lowest-rejected offer. All quantity offered at below the clearing price-point is sold at the clearing price-point times the security’s reference price. Quantity offered at exactly

the clearing price-point is rationed to balance supply and demand, using the proportionate rationing rule.

- *Simultaneous descending clock auction (“clock auction”)*. The securities are auctioned simultaneously. There are two descending clocks, one for high-quality securities and one for low-quality securities, indicating the tentative price-point of each pool. Bidders express the quantities they wish to supply at the current price-points. The price-point is decremented for each pool of securities that has excess supply, and bidders again express the quantities they wish to supply at the new price-points. This process repeats until supply is made equal to demand. The tentative price-points and assignments then become final. Details of the design are presented in Ausubel and Cramton (2008), which you received earlier.

The proportionate rationing rule only plays a role in the event that multiple bidders make reductions at the clearing price. The rule then accepts the reductions at the clearing price in proportion to the size of each bidder’s reduction at the clearing price. Thus, if a reduction of 300 is needed to clear the market and two bidders made reductions of 400 and 200 at the clearing price, then the reductions are rationed proportionately: the first is reduced by 200 and the second is reduced by 100. The actual reduction of the first bidder is twice as large as the second bidder, since the first bidder’s reduction as bid is twice as large as the second bidder’s reduction. Ties can generally be avoided by refraining from bidding price-points that are round numbers, instead specifying price-points to odd one-hundredths of a percent (e.g., 98.42).

The clock auction has an activity rule to encourage price discovery. In particular, for each security pool, the quantities bid must be an upward-sloping supply curves as expressed in activity points. More precisely, *activity points*—equal to the reference price \times quantity, summed over the four securities in the pool—are computed for each bid. The number of activity points is not permitted to increase as the price-point descends. Thus, you are allowed to switch quantities across the four securities in a pool, but your total activity points for the pool cannot increase as the auction progresses.

The same activity rule applies to the sealed-bid auction, but then in a single round of bidding.

In each session, you will participate in four different auctions:

1. Sealed-bid auction, with more precise reference prices
2. Clock auction, with more precise reference prices
3. Sealed-bid auction, with less precise reference prices
4. Clock auction, with less precise reference prices

On Tuesday and Thursday, the order of auctions will be as above. On Wednesday and Friday, the auctions with less precise reference prices will be done first.

In each session, one of your two auctions with more precise reference prices and one of your two auctions with less precise reference prices will be selected at random. Your take-home payment from the session will be based on the number of experimental dollars that you earn in these two auctions only.

1 Securities

In each auction, the Treasury has a demand for each pool of securities: high quality and low quality. The bidders have bidder-specific private information about security values and liquidity needs.

In each session, two sets of bidder-specific private information are drawn independently from the same probability distributions. The first set is used in the auctions with more precise reference prices (auctions 1 and 2); the second set is used in the auctions with less precise reference prices (auctions 3 and 4). You are given no feedback following each sealed-bid auction; thus, your behavior in the subsequent clock auction cannot be influenced by the outcome in the prior sealed-bid auction.

The bidders differ in their security holdings as shown in Table 1.

Table 1. Holdings of securities by bidder and security in thousand \$ of face value

	High-Quality Securities				Low-Quality Securities				Total
	H1	H2	H3	H4	L1	L2	L3	L4	
Bidder 1	20,000					5,000	5,000	10,000	40,000
Bidder 2		20,000			10,000		5,000	5,000	40,000
Bidder 3			20,000		5,000	10,000		5,000	40,000
Bidder 4				20,000	5,000	5,000	10,000		40,000
Bidder 5		5,000	5,000	10,000	20,000				40,000
Bidder 6	10,000		5,000	5,000		20,000			40,000
Bidder 7	5,000	10,000		5,000			20,000		40,000
Bidder 8	5,000	5,000	10,000					20,000	40,000
Total	40,000	40,000	40,000	40,000	40,000	40,000	40,000	40,000	
Expected price	75	75	75	75	25	25	25	25	
Expected value	30,000	30,000	30,000	30,000	10,000	10,000	10,000	10,000	
Total value	120,000				40,000				

Thus, there are four holders of each security: one large (50%), one medium (25%), and two small (12.5% each). Each bidder holds 20,000 (\$k of face value) of high-quality securities and 20,000 (\$k of face value) of low-quality securities.

The four high-quality securities are pooled together and sold as a pool; the four low-quality securities are pooled together and sold as a second pool. Whether done as a sealed-bid auction or done as a clock auction, the two pools are auctioned simultaneously. The Treasury has a budget of \$30,000k for high-quality securities and a budget of \$10,000k for low-quality securities. Thus, given the expected prices of 75 cents on the dollar for high-quality and 25 cents on the dollar for low-quality (see below), the Treasury can be expected to buy a quantity of about 40,000 (\$k of face value), or 25% of face value, of each security pool. Between pools, there is no explicit interaction. However, the bidder's liquidity needs are based on sales from both pools together.

You are one of the eight bidders. You will have the same bidder number in auctions with more precise reference prices (auctions 1 and 2); you will have the same bidder number in auctions with less precise reference prices (auctions 3 and 4). Therefore, your holdings of securities and your signals will be the same in auctions 1 and 2, and they will also be the same in auctions 3 and 4. However, you will have different holdings of securities and signals in auctions 3 and 4, as compared with auctions 1 and 2.

2 Your preferences

From each auction, your payoff depends on your profits and how well your liquidity needs are met.

Common Value Auction

The value of each high-quality security $s \in \{H1, H2, H3, H4\}$ in cents on the dollar is the average of n iid random variables uniformly distributed between 50 and 100:

$$v_s = \frac{1}{n} \sum_{j=1}^n u_{js}, \text{ where } u_{js} \stackrel{iid}{\sim} U[50,100].$$

The value of each low-quality security $s \in \{L1,L2,L3,L4\}$ in cents on the dollar is the average of n iid random variables uniformly distributed between 0 and 50:

$$v_s = \frac{1}{n} \sum_{j=1}^n u_{js}, \text{ where } u_{js} \stackrel{iid}{\sim} U[0,50].$$

For auctions with more precise reference prices, $n = 16$; for auctions with less precise reference prices $n = 12$. The reference price r_s for security s is given by

$$r_s = \frac{1}{n-8} \sum_{j=9}^n u_{js}.$$

Thus, the reference price is based on eight realizations in the more precise case (1/2 of all realizations) and on four realizations in the less precise case (1/3 of all realizations). Reference prices are made public before each auction starts.

For each 5,000 of security holdings, bidder i receives as private information one of the realizations u_{js} . Thus, bidder 1, who holds 20,000 of security 1, gets four realizations. In this way, those with larger holdings have more precise information about the security's value. Observe that this specification requires the holders of each given security to receive collectively a total of eight realizations. Since there are eight realizations available (besides the ones that form the reference price), each of the realizations u_{js} ($i = 1, \dots, 8$) can be observed by exactly one bidder.

Suppose that the auction clearing price-point is p_H for the high-quality pool and p_L for the low-quality pool, where the price-point in the auction is stated as a fraction of the reference price. Then $p_s = p_H r_s$ for $s \in \{H1,H2,H3,H4\}$ and $p_s = p_L r_s$ for $s \in \{L1,L2,L3,L4\}$.

If a bidder sells the quantity q_s (in thousand \$ of face value) of the security s at the price p_s , then profit (in thousand \$) is

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_s (p_s - v_s) q_{is}.$$

The 1/100 factor in the formula above and other formulas involving price is to convert cents into dollars.

Liquidity

Each bidder has a liquidity need, L_i in thousands, which is drawn iid from the uniform distribution on the interval [2500, 7500]. Each bidder knows his own liquidity need, but not that of the other bidders. The bidder receives a bonus of \$1 for every dollar of sales to the Treasury up to his liquidity need:

$$\min \left[L_i, \frac{1}{100} \sum_s p_s q_{is} \right].$$

Payoff of bidder from an auction

Combining the profit and the liquidity penalty results in the bidder's total payoff

$$U_i(p, q_i, v) = \begin{cases} \frac{1}{100} \sum_s (2p_s - v_s) q_{is} & \text{if } \frac{1}{100} \sum_s p_s q_{is} < L_i \\ L_i + \frac{1}{100} \sum_s (p_s - v_s) q_{is} & \text{otherwise} \end{cases}$$

Thus, an additional dollar of cash is worth two dollars when the bidder's liquidity need is not satisfied, but is worth one dollar when the liquidity need is satisfied.

3 Bidder tool and auction system

You have an Excel workbook that contains your private information for each auction. The tool has features that will help your decision making in each of the auctions. Each auction has its own sheet in the tool. It is essential that you are working from the correct sheet for each auction. For example the sealed bid auction with more precise reference prices is the sheet named *Sealed Bid More*.

Bidding is done via a commercial auction system customized to this setting. You use the web browser to connect to the auction system. For each auction, you must go to a new auction site. The URL for the auction site is given in the bidder tool on the particular auction sheet. Once at the correct auction site, log in with the user name and password given on your auction sheet.

The auction system is easy to use. You will have an opportunity to use it in the training seminar.

An important feature of the tool is the calculation of expected security values conditional on information you specify. In a common value auction, it is important for you to condition your estimate of value on your signal *and the information winning conveys*. Since your bid is only relevant in the event that you win, you should set your bid to maximize your payoff in that event. In this way, you avoid the winner's curse, which in this case is the tendency of a naïve bidder to lose money by selling shares at prices below what they are worth. In addition, in the clock auctions, the bidder also must condition on any information revealed through the bidding process.

The bidding tool provides one flexible method of calculating an appropriate conditional expected value for each security. In particular, the tool is set up to calculate the conditional expected values for each security, given the information that you know—the reference price and your own signals—and your best guesses for the relevant other signals. Making appropriate guesses for the other signals is an important element of your strategy. Once these guesses are made, the tool will calculate the common value of the security based on your entries. In the clock auction, you can adjust your estimates of other signals as you learn from the quantity drops of the other bidders.

The tool also lets you keep track of your liquidity bonus based on estimates that you enter for expected prices and expected quantities sold of each security.

Further details of the tool will be explained in the training seminar.

4 Bidding strategy

The auction environment has five complicating features:

- *Common value auction*. You have an imperfect estimate of each security's common value.
- *Divisible good auction*. Your bid is a supply curve, specifying the quantity you wish to sell at various price-points.

- *Demand for pool of securities.* The Treasury does not have a demand for individual securities, but for pools of securities (high- and low-quality pools).
- *Asymmetric holdings.* Bidders have different holdings of securities.
- *Liquidity need.* You have a specific liquidity need that is met through selling shares from your portfolio of eight securities.

The combination of these factors makes a complete equilibrium analysis difficult or impossible, even when we make strong simplifying assumptions. For this reason we will refrain from providing any sort of benchmark strategy.

Your challenge is to determine your own strategy to maximize your payoff that reflects all aspects of the auction environment. The best response in this auction, as in any auction is a best response to what the other bidders are actually doing.

It will be helpful to have an appreciation for the probability density of the common value in various circumstances.

Figure 1 displays the pdf of the common value for low-quality securities in the more precise case by the size of the bidder's holdings, assuming that all the known signals take on the mean value of 25. Thus, when the bidder holds 20,000 of the security there are 4 unknown signals and the standard deviation is 1.8; when the bidder holds 10,000 there are 6 unknown signals and the standard deviation is 2.2; when the bidder holds 5,000 there are 7 unknown signals and the standard deviation is 2.5.

Figure 1. Probability density of common value in more precise case by size of holdings





Figure 2 displays the pdf of the common value for low-quality securities in the less precise case by the size of the bidder's holdings, assuming that all the known signals take on the mean value of 25. Thus, when the bidder holds 20,000 of the security there are 4 unknown signals and the standard deviation is 2.4; when the bidder holds 10,000 there are 6 unknown signals and the standard deviation is 2.9; when the bidder holds 5,000 there are 7 unknown signals and the standard deviation is 3.2.

Figure 2. Probability density of common value in less precise case by size of holdings



If you have any questions during the experiment, please raise your hand and an assistant will help you.

Remember your overall goal is to maximize your experimental payoff in each auction. You should think carefully about what strategy is best apt to achieve this goal.

Many thanks for your participation.

Appendix D

A Common-Value Auction: Bidder Instructions for Experiment 3.1

5 November 2008

In this experiment, you are a bidder in a series of auctions conducted over two sessions (Thursday, 6 November and Saturday, 8 November). Each session is held on a different day and consists of four different auctions. Bidders are randomly assigned to each auction. Thus, you do not know who the other bidders are in any of your auctions. You will be bidding from your private cubical, which includes a computer and a bidder tool that is unique to you. We ask that you refrain from talking during the experiment. If you need assistance at any time, just raise your hand and one of the experimental staff will assist you.

You will be paid at the end of the experiment. In each session, your earnings will be based on your payoff from two randomly selected auctions (out of four total auctions). Your take-home payment is then proportional to your total experimental earnings from sessions 1 and 2. In particular, you will receive *\$0.40 in take-home pay for every one-million experimental dollars that you earn*. Throughout the remainder of the document, dollar figures refer to your experimental payoff unless explicitly stated otherwise—and the “millions” are generally suppressed. When explicitly stated, your real dollar payment will be referred to as your *take-home payment*. We anticipate that each of you will earn a take-home payment of about \$100 per experimental session on average. However, your actual take-home payment will depend on your bidding strategy, the bidding strategies of the other bidders you face, and the particular realizations of numerous random events.

In each auction, you will compete with other bidders to sell your holdings of eight securities to the Treasury. The eight different securities have different values. However, for each security, the bidders have a common value, which is unknown. Each bidder has an estimate of the common value. You profit by selling securities to the Treasury at prices above the securities’ true values. Unlike in previous experiments, you *do not* value liquidity. Thus your payoffs are based solely on profits from your sale of securities to the Treasury.

Two formats are used:

- *Simultaneous uniform-price sealed-bid auction (“sealed-bid auction”)*. Bidders simultaneously submit supply curves for each of the eight securities. Supply curves are non-decreasing (i.e. upward-sloping) step functions. The auctioneer then forms the aggregate supply curve and crosses it with the Treasury’s demand. The clearing price is the lowest-rejected offer. All quantity offered below the clearing price is sold at the clearing price. Quantity offered at the clearing price is rationed to balance supply and demand, using the proportionate rationing rule.
- *Simultaneous descending clock auction (“clock auction”)*. The securities are auctioned simultaneously. There is a price “clock” for each security indicating its tentative price per unit of quantity. Bidders express the quantities they wish to supply at the current prices. The price is decremented for each security that has excess supply, and bidders again express the quantities they wish to supply at the new prices. This process repeats until supply is made equal to demand. The tentative prices and assignments then become final. Details of the design are presented in Ausubel and Cramton (2008), which you received earlier.

The proportionate rationing rule only plays a role in the event that multiple bidders make reductions at the clearing price. The rule then accepts the reductions at the clearing price in proportion to the size of each bidder's reduction at the clearing price. Thus, if a reduction of 300 is needed to clear the market and two bidders made reductions of 400 and 200 at the clearing price, then the reductions are rationed proportionately: the first is reduced by 200 and the second is reduced by 100. The actual reduction of the first bidder is twice as large as the second bidder, since the first bidder's reduction as bid is twice as large as the second bidder's reduction. Ties can generally be avoided by refraining from bidding price-points that are round numbers, instead specifying price-points to odd one-hundredths of a percent (e.g., 98.42).

The clock auction has an activity rule to encourage price discovery. In particular, for each security, the quantities bid must be an upward-sloping supply curve. The quantity of a security is not permitted to increase as the price descends.

In each session, you will participate in *two pairs* of auctions in the following order:

1. 4-bidder sealed-bid, first pair. A sealed bid auction with four bidders.
2. 4-bidder clock, first pair. A clock auction with four bidders. The values of securities and your signals will be identical to those in the sealed-bid above.
3. 4-bidder sealed-bid, second pair. A sealed bid auction with four bidders.
4. 4-bidder clock, second pair. A clock auction with four bidders. The values of securities and your signals will be identical to those in the sealed-bid above.

Each session, one of your first pair of auctions and one of your second pair of auctions will be selected at random. Your take-home payment from the session will be based on the number of (million) laboratory dollars that you earn in these two auctions only.

1 Securities

In each auction, eight securities are purchased by the Treasury. The bidders are symmetric, before the draws of bidder-specific private information about security values.

In each session, two sets of bidder-specific private information are drawn independently from the same probability distributions. The first set is used in the first pair of auctions (auctions 1 and 2); the second set is used in the second pair of auctions (auctions 3 and 4). You are given no feedback following the sealed-bid auctions; thus, your behavior in the clock auctions cannot be influenced by outcomes in the sealed-bid auctions of a session.

In each auction, the Treasury demand is 1000 shares of each security, where each share corresponds to \$1 million of face value. Each bidder has holdings of 1000 shares of each security. Thus, it is possible for a single bidder to fully satisfy the Treasury demand for a particular security; that is, for each security there may be just a single winner or there may be multiple winners. One quarter of the total available shares will be purchased by the Treasury.

2 Your preferences

From each auction, your payoff depends on your profits from the sale of securities to the Treasury.

Common Value Auction

The value of each security in cents on the dollar is the average of eight iid random variables uniformly distributed between 0 and 100:

$$v_s = \frac{1}{8} \sum_{i=1}^8 u_{is}, \text{ where } u_{is} \stackrel{iid}{\sim} U[0,100].$$

Suppose you are bidder i .

Your private information about security s is the realization u_{is} . Notice that the common value is the average of eight uniform draws, so that only the first four draws are revealed (as there are only four bidders in the auction).

If you sell the quantity q_s of the security s at the price p_s , then your profit (in million \$) is

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_{s=1}^8 (p_s - v_s) q_{is}.$$

3 Bidder tool and auction system

You have an Excel workbook that contains your private information for each auction. The tool has features that will help your decision making in each of the auctions. Each auction has its own sheet in the tool. It is essential that you are working from the correct sheet for each auction. For example the sealed-bid, first-pair auction is the sheet named *Sealed Bid First Pair*.

For each of the auctions, bidding is done via a commercial auction system customized to this setting. You use the web browser to connect to the auction system. For each auction, you must go to a new auction site. The URL for the auction site is given in the bidder tool on the particular auction sheet, Sealed Bid First Pair, Clock First Pair, Sealed Bid Second Pair, or Clock Second Pair. Once at the correct auction site, log in with the user name and password given on your auction sheet.

The auction system is easy to use. It is identical to the system you used for bidding in all previous experiments.

An important feature of the tool is the calculation of expected security values conditional on information you specify. In a common value auction, it is important for you to condition your estimate of value on your signal *and the information winning conveys*. Since your bid is only relevant in the event that you win, you should set your bid to maximize your payoff in that event. In this way, you avoid the winner's curse, which in this case is the tendency of a naïve bidder to lose money by selling shares at prices below what they are worth. In addition, in the clock auctions, the bidder also must condition on any information revealed through the bidding process. The tool provides one flexible method of calculating an appropriate conditional expected value for each security.

4 Bidding strategy

The auction environment has two complicating features:

- *Common value auction*. You have an imperfect estimate of the good's common value.
- *Divisible good auction*. Your bid is a supply curve, specifying the quantity you wish to sell at various prices.

The combination of these factors makes a complete equilibrium analysis difficult. Nonetheless, equilibrium analysis is possible in a simplified environment. To aid your thinking about strategy, we

discuss a particular strategy, which abstracts from the complications of a divisible good auction. In particular, we assume that each bidder submits a flat supply schedule; that is, the bidder offers to sell all of her holdings of a particular security at a specified price.

With these assumptions it is possible to calculate an equilibrium strategy, which we call the *benchmark strategy*. The analysis of this strategy will be helpful in thinking about the common value feature of the auction environment.

We wish to emphasize that the benchmark strategy focuses on only one element of the auction. Your challenge is to determine your own strategy to maximize your payoff that reflects all aspects of the auction environment.


4.1 Common value distribution

As mentioned earlier, the value of each security in cents on the dollar is the average of eight iid random variables uniformly distributed between 0 and 100:

$$v_s = \frac{1}{8} \sum_{i=1}^8 u_{is}, \text{ where } u_{is} \stackrel{iid}{\sim} U[0,100].$$

The pdf and cdf of the common value are shown in Figure 1. The common value has a mean of 50 and a standard deviation of 10.2. Notice that the common value is approximately normally distributed, since it is the sum of many independent draws.

Figure 1. Probability density and cumulative distribution of common value



You are bidder i . Your private information about security s is the realization u_{is} .

4.2 Sealed-bid uniform-price reverse auction

Under our strong simplifying assumptions, the auction is equivalent to a unit-supply common value auction with uniform pricing. In this case, just as in a single-item second-price auction, your bid does not determine what you pay, only the likelihood of winning, thus the best strategy is to bid your true conditional expected value for the good. The trick, however, is to condition on your signal and the information that winning conveys.

In the 4-bidder auction, under the benchmark assumptions, there is only a single winner, so the correct condition is as derived in Milgrom and Weber (1982). Your bid is your expected value conditional on your signal being the lowest and on the second-lowest signal being equal to yours:

$$b_{is} = E[v_s | u_1 = u_{is}, u_2 = u_1],$$

where u_1 is the lowest signal and u_2 is the second-lowest signal. The reason you condition on the second-lowest signal being as low as yours is that the bid must be optimal when it is on the margin and thus impacts whether you win. Your bid becomes marginal and hence decisive only in the event that you tie with the second-lowest.

Since the signals are uniformly distributed, it is easy to calculate the above conditional values. In both case, the conditional value is a linear function of your signal.

In the 4-bidder auction, there is a single winner. The conditioning is that you win, the next lowest bidder has your value, and the remaining two bidders are evenly distributed above your value:

$$\text{4-bidder sealed-bid strategy: } b_{is} = \frac{1}{8} \left(2u_{is} + 2 \left(\frac{u_{is} + 100}{2} \right) + 4 \cdot 50 \right) = \frac{3}{8} u_{is} + \frac{75}{2}.$$

4.3 Descending clock auction

In the clock auction, under the benchmark assumption, you will observe the price at which other bidders drop out. This provides additional information on which to condition your bid. Here we assume that the price clock is continuous. In the actual experiment, the price clock is discrete, and although bidders can make reductions at any price, you will only learn the aggregate supply at the end of round price. You may want to assume the quantity reduction occurred half-way between the prior price and the ending price.

When the clock starts, you calculate your dropout point in the same way as above. As the price clock falls, one of the other bidders may drop out. When the first bidder drops out, you can calculate this bidder's draw from the following equation.

$$u_{8s} = \frac{P_{\text{dropout}} - \text{intercept}}{\text{slope}}$$

where the slope and intercept are taken the formulas above. With this new information on which to condition your bid, the revised optimal bid for the 4-bidder clock auction is straightforward to calculate.

4-bidder clock strategy

$$\text{No one has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + 2 \left(\frac{u_{is} + 100}{2} \right) + 4 \cdot 50 \right)$$

$$\text{One bidder has dropped out: } b_{is} = \frac{1}{8} \left(2u_{is} + u_4 + 1 \left(\frac{u_{is} + u_4}{2} \right) + 4 \cdot 50 \right).$$

$$\text{Two bidders have dropped out: } b_{is} = \frac{1}{8} (2u_{is} + u_4 + u_3 + 4 \cdot 50).$$

Note that the above formulas are all linear in the dropout price, so it is easy to invert to get compute the bidder's signal.

4.4 Moving beyond the benchmark strategy

The bidding tool is set up to calculate the conditional expected values assuming the benchmark strategy. Of course, you (and others) may well deviate from the benchmark strategy.

The bidding tool allows for this variation in a number of ways. First, in the sealed-bid auctions your bid can be any upward sloping step function to account for possible supply reduction by you or others. Second, in the clock auctions, the tool lets you interpret what a dropout is and where it occurs. This is useful and necessary when bidders make partial reductions of supply. In addition, although the tool will calculate a particular signal based on a dropout, you are free to type in any signal inference you like. Whatever you type as “my best guess” will be used in the calculation of the conditional expected value.

If you have any questions during the experiment, please raise your hand and an assistant will help you.

Remember your overall goal is to maximize your experimental payoff in each auction. You should think carefully about what strategy is best apt to achieve this goal.

Many thanks for your participation.

Appendix E

A Common-Value Reference-Price Auction: Bidder Instructions for Experiment 3.2

8 November 2008

In this experiment, you are a bidder in a series of auctions conducted over two sessions (Monday, 10 November and Tuesday, 11 November). Each session is held on a different day and consists of four different auctions. Bidders are randomly assigned to each auction. Thus, you do not know who the other bidders are in any of your auctions. You will be bidding from your private cubical, which includes a computer and a bidder tool that is unique to you. We ask that you refrain from talking during the experiment. If you need assistance at any time, just raise your hand and one of the experimental staff will assist you.

You will be paid at the end of the experiment. In each session, your earnings will be based on your payoff from two randomly selected auctions (out of four total auctions). Your take-home payment is then proportional to your total experimental earnings from sessions 1 and 2. In particular, you will receive *\$0.30 in take-home pay for every one thousand experimental dollars (\$k) that you earn*. Throughout, dollar figures refer to your experimental payoff unless explicitly stated otherwise—and the “thousands” are generally suppressed. When explicitly stated, your real dollar payment will be referred to as your *take-home payment*. We anticipate that each of you will earn a take-home payment of about \$100 per experimental session on average. However, your actual take-home payment will depend on your bidding strategy, the bidding strategies of the other bidders you face, and the particular realizations of numerous random events.

In each auction, you will compete with other bidders to sell your holdings of eight securities to the Treasury. The eight securities are split into two pools: four securities are low quality and four are high quality. The eight different securities have different values. However, for each security, the bidders have a common value, which is unknown. Each bidder has an estimate of the common value. You profit by selling securities to the Treasury at prices above the securities’ true values. The Treasury has allocated a particular budget for the purchase of each pool of securities within each auction. Its demand for high-quality securities is distinct from its demand for low-quality securities.

Before each auction, the auctioneer assigns each security a *reference price* (expressed in cents on the dollar of face value), which represents the Treasury’s best estimate of what each security is worth. For example, a high-quality security might be assessed to be worth 75 cents on the dollar, while a low-quality security might be assessed to be worth 25 cents on the dollar. A *reference-price auction* determines the price-point—a percentage of the reference price—for each pool of securities. A winning bidder is paid the pool’s price-point \times the security’s reference price for each unit of the security sold.

Two formats are used:

- *Simultaneous uniform-price sealed-bid auction (“sealed-bid auction”)*. Bidders simultaneously submit supply curves for each of the eight securities. Supply curves are non-decreasing (i.e. upward-sloping) step functions, offering a quantity at each price-point. The auctioneer then forms the aggregate supply curve and crosses it with the Treasury’s demand. This is done separately for each pool (i.e., for high- and low-quality securities, separately). The clearing price-point is the lowest-rejected offer. All quantity offered at below the clearing price-point is sold at the clearing price-point times the security’s reference price. Quantity offered at exactly

the clearing price-point is rationed to balance supply and demand, using the proportionate rationing rule.

- *Simultaneous descending clock auction (“clock auction”)*. The securities are auctioned simultaneously. There are two descending clocks, one for high-quality securities and one for low-quality securities, indicating the tentative price-point of each pool. Bidders express the quantities they wish to supply at the current price-points. The price-point is decremented for each pool of securities that has excess supply, and bidders again express the quantities they wish to supply at the new price-points. This process repeats until supply is made equal to demand. The tentative price-points and assignments then become final. Details of the design are presented in Ausubel and Cramton (2008), which you received earlier.

The proportionate rationing rule only plays a role in the event that multiple bidders make reductions at the clearing price. The rule then accepts the reductions at the clearing price in proportion to the size of each bidder’s reduction at the clearing price. Thus, if a reduction of 300 is needed to clear the market and two bidders made reductions of 400 and 200 at the clearing price, then the reductions are rationed proportionately: the first is reduced by 200 and the second is reduced by 100. The actual reduction of the first bidder is twice as large as the second bidder, since the first bidder’s reduction as bid is twice as large as the second bidder’s reduction. Ties can generally be avoided by refraining from bidding price-points that are round numbers, instead specifying price-points to odd one-hundredths of a percent (e.g., 98.42).

The clock auction has an activity rule to encourage price discovery. In particular, for each security pool, the quantities bid must be an upward-sloping supply curves as expressed in activity points. More precisely, *activity points*—equal to the reference price \times quantity, summed over the four securities in the pool—are computed for each bid. The number of activity points is not permitted to increase as the price-point descends. Thus, you are allowed to switch quantities across the four securities in a pool, but your total activity points for the pool cannot increase as the auction progresses.

The same activity rule applies to the sealed-bid auction, but then in a single round of bidding.

In each session, you will participate in *two pairs* of auctions in the following order:

5. Sealed-bid auction, first pair.
6. Clock auction, first pair.
7. Sealed-bid auction, second pair.
8. Clock auction, second pair.

In all cases, the reference prices will be based on 4 signals. That is, the reference prices will be analogous to those used in the *less precise* auctions from experiment 2.

Each session, one of your first pair of auctions and one of your second pair of auctions will be selected at random. Your take-home payment from the session will be based on the number of (million) laboratory dollars that you earn in these two auctions only.

1 Securities

In each auction, the Treasury has a demand for each pool of securities: high quality and low quality. The bidders have bidder-specific private information about security values.

In each session, two sets of bidder-specific private information are drawn independently from the same probability distributions. The first set is used in the first pair of auctions (auctions 1 and 2); the second set is used in the second pair of auctions (auctions 3 and 4). You are given no feedback following each sealed-bid auction; thus, your behavior in the subsequent clock auction cannot be influenced by the outcome in the prior sealed-bid auction.

The bidders differ in their security holdings as shown in Table 1.

Table 1. Holdings of securities by bidder and security in thousand \$ of face value

	High-Quality Securities				Low-Quality Securities				Total
	H1	H2	H3	H4	L1	L2	L3	L4	
Bidder 1	20,000					5,000	5,000	10,000	40,000
Bidder 2		20,000			10,000		5,000	5,000	40,000
Bidder 3			20,000		5,000	10,000		5,000	40,000
Bidder 4				20,000	5,000	5,000	10,000		40,000
Bidder 5		5,000	5,000	10,000	20,000				40,000
Bidder 6	10,000		5,000	5,000		20,000			40,000
Bidder 7	5,000	10,000		5,000			20,000		40,000
Bidder 8	5,000	5,000	10,000					20,000	40,000
Total	40,000	40,000	40,000	40,000	40,000	40,000	40,000	40,000	
Expected price	75	75	75	75	25	25	25	25	
Expected value	30,000	30,000	30,000	30,000	10,000	10,000	10,000	10,000	
Total value	120,000				40,000				

Thus, there are four holders of each security: one large (50%), one medium (25%), and two small (12.5% each). Each bidder holds 20,000 (\$k of face value) of high-quality securities and 20,000 (\$k of face value) of low-quality securities.

The four high-quality securities are pooled together and sold as a pool; the four low-quality securities are pooled together and sold as a second pool. Whether done as a sealed-bid auction or done as a clock auction, the two pools are auctioned simultaneously. The Treasury has a budget of \$30,000k for high-quality securities and a budget of \$10,000k for low-quality securities. Thus, given the expected prices of 75 cents on the dollar for high-quality and 25 cents on the dollar for low-quality (see below), the Treasury can be expected to buy a quantity of about 40,000 (\$k of face value), or 25% of face value, of each security pool. Between pools, there is no explicit interaction.

You are one of the eight bidders. You will have the same bidder number in the first pair of auctions (auctions 1 and 2); you will have the same bidder number in the second pair of auctions (auctions 3 and 4). Therefore, your holdings of securities and your signals will be the same in auctions 1 and 2, and they will also be the same in auctions 3 and 4. However, you will have different holdings of securities and signals in auctions 3 and 4, as compared with auctions 1 and 2.

2 Your preferences

From each auction, your payoff depends on your profits.

Common Value Auction

The value of each high-quality security $s \in \{H1,H2,H3,H4\}$ in cents on the dollar is the average of 12 iid random variables uniformly distributed between 50 and 100:

$$v_s = \frac{1}{12} \sum_{j=1}^{12} u_{js}, \text{ where } u_{js} \stackrel{iid}{\square} U[50,100].$$

The value of each low-quality security $s \in \{L1, L2, L3, L4\}$ in cents on the dollar is the average of 12 iid random variables uniformly distributed between 0 and 50:

$$v_s = \frac{1}{12} \sum_{j=1}^{12} u_{js}, \text{ where } u_{js} \stackrel{iid}{\sim} U[0, 50].$$

The reference price r_s for security s is given by

$$r_s = \frac{1}{4} \sum_{j=9}^{12} u_{js}.$$

Thus, the reference price is based on four realizations (1/3 of all realizations). Reference prices are made public before each auction starts.

For each 5,000 of security holdings, bidder i receives as private information one of the realizations u_{js} . Thus, bidder 1, who holds 20,000 of security 1, gets four realizations. In this way, those with larger holdings have more precise information about the security's value. Observe that this specification requires the holders of each given security to receive collectively a total of eight realizations. Since there are eight realizations available (besides the ones that form the reference price), each of the realizations u_{js} ($i = 1, \dots, 8$) can be observed by exactly one bidder.

Suppose that the auction clearing price-point is p_H for the high-quality pool and p_L for the low-quality pool, where the price-point in the auction is stated as a fraction of the reference price. Then $p_s = p_H r_s$ for $s \in \{H1, H2, H3, H4\}$ and $p_s = p_L r_s$ for $s \in \{L1, L2, L3, L4\}$.

If a bidder sells the quantity q_s (in thousand \$ of face value) of the security s at the price p_s , then profit (in thousand \$) is

$$\pi_i(p, q_i, v) = \frac{1}{100} \sum_s (p_s - v_s) q_{is}.$$

The 1/100 factor in the formula above and other formulas involving price is to convert cents into dollars.

3 Bidder tool and auction system

You have an Excel workbook that contains your private information for each auction. The tool has features that will help your decision making in each of the auctions. Each auction has its own sheet in the tool. It is essential that you are working from the correct sheet for each auction. For example the sealed bid auction from the first pair is the sheet named *Sealed Bid First Pair*.

Bidding is done via a commercial auction system customized to this setting. You use the web browser to connect to the auction system. For each auction, you must go to a new auction site. The URL for the auction site is given in the bidder tool on the particular auction sheet. Once at the correct auction site, log in with the user name and password given on your auction sheet.

The auction system is easy to use. It is identical to the system you used for bidding in all previous experiments.

An important feature of the tool is the calculation of expected security values conditional on information you specify. In a common value auction, it is important for you to condition your estimate of value on your signal *and the information winning conveys*. Since your bid is only relevant in the event that you win, you should set your bid to maximize your payoff in that event. In this way, you avoid the winner's curse, which in this case is the tendency of a naïve bidder to lose money by selling shares at

prices below what they are worth. In addition, in the clock auctions, the bidder also must condition on any information revealed through the bidding process.

The bidding tool provides one flexible method of calculating an appropriate conditional expected value for each security. In particular, the tool is set up to calculate the conditional expected values for each security, given the information that you know—the reference price and your own signals—and your best guesses for the relevant other signals. Making appropriate guesses for the other signals is an important element of your strategy. Once these guesses are made, the tool will calculate the common value of the security based on your entries. In the clock auction, you can adjust your estimates of other signals as you learn from the quantity drops of the other bidders.

4 Bidding strategy

The auction environment has four complicating features:

- *Common value auction.* You have an imperfect estimate of each security's common value.
- *Divisible good auction.* Your bid is a supply curve, specifying the quantity you wish to sell at various price-points.
- *Demand for pool of securities.* The Treasury does not have a demand for individual securities, but for pools of securities (high- and low-quality pools).
- *Asymmetric holdings.* Bidders have different holdings of securities.


The combination of these factors makes a complete equilibrium analysis difficult or impossible, even when we make strong simplifying assumptions. For this reason we will refrain from providing any sort of benchmark strategy.

Your challenge is to determine your own strategy to maximize your payoff that reflects all aspects of the auction environment. The best response in this auction, as in any auction is a best response to what the other bidders are actually doing.

It will be helpful to have an appreciation for the probability density of the common value in various circumstances.

Figure 1 displays the pdf of the common value for low-quality securities by the size of the bidder's holdings, assuming that all the known signals take on the mean value of 25. Thus, when the bidder holds 20,000 of the security there are 4 unknown signals and the standard deviation is 2.4; when the bidder holds 10,000 there are 6 unknown signals and the standard deviation is 2.9; when the bidder holds 5,000 there are 7 unknown signals and the standard deviation is 3.2.

Figure 1. Probability density of common value by size of holdings



If you have any questions during the experiment, please raise your hand and an assistant will help you.

Remember your overall goal is to maximize your experimental payoff in each auction. You should think carefully about what strategy is best apt to achieve this goal.

Many thanks for your participation.