A Reverse Auction for Toxic Assets

Olivier Armantier, Charles A. Holt, and Charles R. Plott*

March 10, 2010

Abstract

The proposed 2008 TARP auction was intended to facilitate rapid purchases by the U.S. Treasury of a wide array of mortgage-backed securities in order to remove these "toxic assets" from the portfolios of financially stressed banks. The Treasury had decided on a Reference Price design whereby bids to sell different securities would be normalized or "scored" by reference prices that reflect estimates of relative values. Although the auction was suspended just prior to the 2008 Presidential election, we continued with a series of laboratory experiments aimed at testing the performance of the Reference Price design relative to other auction formats. The results of the experiments presented in this paper indicate that a simple Reference Price auction can be an effective mechanism for avoiding serious effects of adverse selection and strategic bid manipulation, even when reference prices are not set accurately. An econometric analysis of bidding patterns and auction outcomes reveals how underlying behavior produces efficiency differences in this common-value framework.

Keywords: reverse auction, auction design, common value auction, mortgage backed securities, laboratory experiments

I. Introduction

In the fall of 2008, as the financial crisis reached a critical stage, the U.S. Treasury (hereafter Treasury) decided to conduct a series of auctions in order to purchase diverse mortgage backed securities (hereafter MBS) from financial institutions. With a budget of about $700 billion, this would have been the largest auction in history, surpassing even the largest spectrum auctions by a factor of 10. Since no appropriate "off-the-shelf" design existed for this purpose, a new auction format had to be rapidly developed. The scale of the proposed auction and immediacy of the crisis generated considerable discussion among economists (e.g., Varian, 2008; Ausubel and Cramton, 2008). After consulting with academic experts, the Treasury decided to implement a new design, a

* Federal Reserve Bank of New York (Armantier), University of Virginia (Holt), and Caltech (Plott). We wish to thank Claudia Antonacci, Andrew Barr, Rachel Blank, Anne Dragnova, Sara St. Hilaire, Stephanie Lawrence, Courtney Mallow, Karl Schurter, and Sean Sullivan for research assistance. We also thank seminar and conference participants at the 2009 Workshop on Neuroeconomics and Psychology in Toulouse, the 2009 Public Economic Conference in Lyon, the 2009 ESA meeting in Tucson, the 2010 ESA meeting in Melbourne, and the University of Technology in Sydney. This research was supported in part by the University of Virginia Bankard Fund for Political Economy, and by the CIRANO lab in Montreal, where some of the experiment sessions were conducted. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System.

"Reference Price" auction. It was unclear, however, how this design would perform. In this paper, we report on an experiment aimed at testing the basic Reference Price auction adopted by the Treasury, as well as comparing how alternatives features may provide some protection against bidders possibly exploiting the Treasury's inferior information about security values.

The design of an auction to purchase MBS as part of the Troubled Asset Relief Program (or TARP) presented a number of unique challenges, including:

1. **Dimensionality:** The universe of possible MBS eligible for purchase by the Treasury was extremely large. In particular, within the realm of residential MBS alone, there were more than 23,000 distinct securities.
2. **Heterogeneity**: MBS are highly heterogeneous as they can differ along several observed and unobserved dimensions including their ratings, the vintages of the underlying mortgages, the locations and/or the amounts of the mortgages, the characteristics of the borrowers (e.g. credit ratings), the mortgages delinquency rates, and the originators of the loans. The ratings assigned by credit agencies to MBS were supposed to reflect these differences, but many observers at that point in time believed that these ratings had lost most of their meaning.
3. **Valuation**: At the time the auction was designed, there was virtually no market for most mortgage-backed securities that the Treasury was considering for purchase, and hence no price information. Moreover, the unprecedented rate of commercial and residential foreclosures over the preceding year made it hazardous to rely on standard simulation models to value these securities.
4. **Informational Asymmetry**: Security owners could be assumed to have access to better information about underlying asset values than outside analysts and prospective buyers (e.g., by observing the repayment streams and default rates of the borrowers in the pool of mortgages).
5. **Concentration of Ownership:** A large fraction of the shares of a specific security could belong to a single financial institution, thereby creating a thin market problem. Although the extent of the problem was unknown, it was believed to apply to many of the securities under consideration.
6. **Time Constraints**: The time constraints were threefold. First, the auction had to be designed and tested within a matter of weeks. Second, for any given auction, the bidding process needed to be completed in a short time frame so as to avoid outside influence and possible sellers' regret.[1] Third, once initiated, the program needed to have a significant and immediate impact to restore market confidence.

The Treasury realized that it might not be able to fully address all of these challenges. In particular, it understood that in the presence of asymmetric information, it is often difficult to avoid inefficiencies. The goal then became to design an auction that would effectively mitigate possible adverse consequences. Two simple auction formats were

[1] Given the high market volatility at the time, it was feared that bidding could be influenced by sudden market developments if the bidding process was to last for more than a few minutes. Conversely, concerns were expressed about the possibility that financial markets could be affected or manipulated during the auction process.

rapidly abandoned. The first was a "Grand Auction," whereby different securities would be pooled together in a single reverse auction and purchased at a single price, irrespective of the type of security accepted. A perceived advantage of this approach was that it could promote competition by letting owners of different securities compete against each other in a single auction. However, because of heterogeneity and informational asymmetry problems, everyone involved recognized that adverse selection could be severe in a simple combined auction. With low bids likely to come in for the more "toxic" assets, it was believed that the Treasury would end up purchasing the lowest-quality securities. This outcome would not necessarily be a problem unless it involved large overpayments by the Treasury, which was likely to be the case in a simple combined auction. The second format to be ruled out was a security-by-security auction. This approach would effectively deal with the unobserved value heterogeneity, as a separate auction would be conducted for each of the thousands of securities. Because of the problems of dimensionality, ownership concentration, and time constraints, however, this approach was considered not only impractical but also undesirable.

In essence, the design the Treasury decided to implement is a version of what is known in the literature as a "scoring auction," in which bids with heterogeneous characteristics are homogenized through the use of a score that incorporates both price and quality or other non-price dimensions of the bids.[2] As further explained in section III, the non-price dimension for the TARP auction was to be based on "reference price" estimates of the relative values of the mortgage backed securities. It was believed that the main benefit of this design is that it could deal with the dimensionality problem while promoting competition and providing some protection against adverse selection. In addition, although undoubtedly sub-optimal in the formal mechanism design sense, it was deemed simple, transparent, and robust enough to be implemented rapidly and efficiently. The Treasury, however, realized that the performance of this design relied heavily on its ability to precisely estimate reference prices. As such an exercise could not be conducted without making mistakes, concerns were expressed about the possibility that bidders could exploit these mistakes at the Treasury's expense. As further discussed in Section III several remedies were proposed to mitigate such a problem. In particular, it was suggested to reveal the Treasury's reference prices only after the auction so that the bidders would not be in a position to observe the Treasury's mistakes at the time they submit their bids.

[2] For instance, for its construction projects, the European Spatial Agency asks bidders to submit proposals consisting of a price and a quality. The bids are then ranked according to a score equal to the ratio of quality to bid price (Armantier, Florens and Richard 1998, Armantier 1999). A similar approach was implemented in the Victoria Bush Tender Auction, in which a government agency first estimated a "biodiversity index" for land management conservation actions that were proposed by participating farmers. The auction was cleared by accepting bids with the highest biodiversity-per-dollar ratios, and then moving down the list until the limited purchase budget was exhausted (Stoneham, et al. 2002). See Cason, Gangadharan, and Duke (2003) for a laboratory study in which the bids were scored by a measure of environmental benefit. Scoring auctions are used in a majority of states for highway construction contracts, with the score being a weighted average of the bid price and the road user cost of time delay. This linear scoring method is known as "A + B bidding." See Asker and Cantillon (2008) for a theoretical analysis of linear scoring auctions and for references to applications that range from electricity reserve supply to public procurement in the European Union.

Given the complexity of the design, the relevant issues could not be addressed directly with economic theory. A commonly used procedure in such situations is to conduct a laboratory experiment that simulates the environment with financially motivated human subjects.[3] The objective of the experiment conducted in this paper is threefold. First, we want to test whether a Reference Price auction performs better (from the Treasury's perspective) than a Grand Auction when reference prices are set accurately by the Treasury. Second, we want to explore how sensitive the Reference Price auction is to reference prices that are incorrectly set by the Treasury. More specifically, we want to test whether announcing noisy (instead of accurate) reference prices before the auction generates inefficiencies, as bidders are in a position to exploit mistakes made by the Treasury. Third, we want to test one of the measure proposed to protect the Treasury against the mistakes it would make in setting reference prices. Namely, we test whether the Reference Price auction produces more efficient outcomes when the noisy reference prices are kept secret at the time of bidding instead of being announced before the auction.

The remainder of the paper is structured as follows. Section II describes the extraordinary environment in which the auction was conceived. The Reference Price auction is explained in Section III. Section IV details the experimental procedures used to compare parallel series of laboratory test auctions, using financially motivated bidders. Section V summarizes the main results of these experiments, in terms of efficiency (value of securities purchased relative to government expenditures) and price discovery (auction clearing prices relative to unobserved security values). In Section VI, the causes of performance differences are investigated with an econometric analysis of bidding behavior and auction outcomes. The final section contains a discussion of extensions and alternative applications of auctions for heterogeneous items.

II. Setting the Stage

In early September of 2008, the adverse consequences of the subprime crisis spread and deepened at an alarming pace. On September 7, the Federal Reserve (hereafter the Fed) took over *Fannie Mae* and *Freddie Mac*, two cornerstones in the securitization of mortgages. Half of the remaining U.S. investment banks ceased to exist on September 15, as *Merrill Lynch* was sold to *Bank of America*, while *Lehman Brothers* filed for bankruptcy protection. The *Lehman Brothers* bankruptcy (the largest in U.S. history) set off what had been described as a devastating tidal wave that immediately sliced through entire financial system. Two days later, the Fed loaned $85 billion to *AIG*. In spite of a

[3] See Holt, Shobe, and Smith (2006) for a survey of the use of laboratory experiments to guide public policy initiatives, including auctions. The auction design team for the Regional Greenhouse Gas Initiative (RGGI) relied on an extensive set of experiments in making their recommendations, which were implemented almost immediately (Holt, et al., 2008). Laboratory experiments were also used to test the combinatorial component of the 700 MHZ Federal Communications Commission spectrum auction conducted in early 2008. In fact, the particular "hierarchical package bidding" procedure that was used there had never been implemented outside of a laboratory before it was used in a multi-billion dollar auction (Goeree and Holt, 2008a). A reverse auction for irrigation permits in south Georgia was patterned after a design that had been tested in a series of experiments involving both farmers and students (Cummings, Holt, and Laury, 2004).

record injection of $180 billion in liquidity by central banks around the world, the situation became critical on September 18: interbank credit markets almost entirely collapsed, money market funds experienced an unprecedented run in excess of half a trillion dollars, while the commercial paper market was on verge of dislocation.[4] At this point in time, some observers felt that the financial system was on the edge of a precipice and could completely melt down in a matter of days.

On the evening of September 18, then Treasury Secretary Henry Paulson and Fed Chairman Ben Bernanke met with members of Congress. The atmosphere in the meeting has been described as grave and somber.[5] Secretary Paulson presented a plan aimed at creating a "firewall" to stop the crisis from spreading from the weakest financial institution to the next. Instead of acting on a case-by-case basis, the plan called for a bold, comprehensive approach to attack the root of the crisis. More specifically, Secretary Paulson proposed to purchase $700 billion in mortgage-backed securities that were perceived to clog the overall flow of credit to financial institutions, corporations, and consumers. The rationale was that the virtual collapse of the MBS market had made it extremely difficult to value these securities, and by extension, the financial institutions that owned them. This increased uncertainty was perceived to be the main reason behind the reluctance of market participants to loan and trade with one another. The hope was that removing these illiquid assets from banks' balance sheets would restore confidence in financial markets.

The initial bill proposed by the Treasury, which was less than three pages long, was rejected in the House of Representatives on September 29. The Dow Jones Industrial Average immediately lost $1.2 trillion in market value (1.7 times more than the TARP bill), or 777 points (the largest single-day point drop in its history). Congress then amended the bill, which swelled to more than 400 pages and $850 billion. When the legislation passed on October 3, however, it was unclear how the mortgage backed securities would be acquired from banks and at what price. A reverse ("low bids win") auction had been mentioned, as it was perceived to be an efficient and transparent mechanism for determining prices when markets are not functioning. Such an approach had been successfully adopted in the late 1980's during the "Saving and Loans Crisis" to purchase distressed real estate. As many noted, however, MBS are complex securities that are different from real estate, and doubts were immediately expressed about the effectiveness of such a straightforward reverse auction in the current context. Since there was no ready-to-use procedure that would fit the Treasury needs, a new type of auction needed to be designed.

On Saturday, September 20, a team at the Federal Reserve Bank of New York proposed a basic structure for government purchases of large numbers of mortgage backed securities

---

[4] Commercial paper is used by corporations to meet short term funding needs to finance (e.g.) payroll or rent.

[5] Chairman Bernanke was reported as saying to a member of Congress: "If we don't do this, we may not have an economy on Monday." Another person remarked: "We are literally maybe days away from a complete meltdown of our financial system" (C. Dodd). Another recounted: "When you listened to him describe it, you gulped." (C. Schumer).

through a reverse auction based on reference prices (Armantier and Vickery 2008). In the week that followed, academics with experience in designing auctions from the ground up were contacted. These experts were grouped into four teams that developed separate proposals. The Treasury decided to implement a design submitted by Jacob Goeree and Charles Holt (2008b), and to associate them with Charlie Plott to refine their design and incorporate elements he had suggested.

This basic reference price design, however, had not been tested, and it was not clear how it would perform. In such situations, it is becoming increasingly common to use laboratory experiments to test and refine the design of specific auction procedures.[6] We began the process of setting up experiments in mid-October of 2008, just weeks before the design consultants were told to "take a vacation" (which presumably meant that the auction was on hold, although an official announcement was not ready at that time). On November 12, the Treasury officially abandoned its plan for the TARP auction. Instead, it decided to use the funds allocated by Congress to take an equity position in some banks. Secretary Paulson explained that the situation had changed and that capital injections were now a more effective approach to address the situation faced by the financial system.

III. The Auction Design Adopted

Although it did not have what economists would characterize as a well defined objective function, the Treasury had specific priorities and concerns when designing the new TARP auction. The first version of the bill submitted to Congress stated that the objectives of the purchase program were i) stabilizing financial markets, and ii) protecting the taxpayer. The Treasury realized it needed to strike a delicate balance between overpaying for the securities (which would harm taxpayers) and underpaying (which would harm financial institutions).[7] Overall, the principal objective was to design an auction yielding a price between what was described as the "fire sale price," at which some of these assets were traded at the time, and the "hold to maturity price" which reflects the stream of mortgage revenues a patient investor would receive at maturity. Additional objectives were also considered, e.g., price discovery and provision of bank liquidity, but these were deemed to be second-order priorities.[8]

---

[6] The experimental study that is the most relevant for the Treasury TARP program is that of Cason, Gangadharan, and Duke (2003), who score bids for land conservation on the basis of a measure of environmental benefit. One major difference from the Treasury situation, however, was that the environmental benefit only mattered to the government purchaser, whereas security values matter to both sellers and the government purchaser of mortgage-backed securities.

[7] Observe that, because of the marked-to-market accounting practices in the U.S., underpayments by the Treasury would have hurt not only the banks that sold a given security, but also the banks that still owned shares of that security, as they most likely would have had to mark down the value of the security in their books. Given the fragility of the entire financial system at the time, such write-downs could have had severe consequences.

[8] There was considerable doubt about the ability of an auction to reveal meaningful price information. In particular, with the $700 billion purchase budget being contemplated at that time, the Treasury was likely to substantially shift the demand function for these assets.

We now describe the basic features of the auction design the Treasury adopted, and how it was supposed to unfold.[9] A series of 20 to 30 different auctions was expected to be conducted, each at a different date. A few days before a given auction, Treasury officials would ask eligible bidders to indicate how many shares of each of a pre-announced set of (say 1,000) securities they would be willing to sell. For each security, a bidder would not be allowed to submit a bid to sell a quantity that exceeded what was listed in this "nomination" phase. In addition, the Treasury was considering asking participating bidders not to buy or sell any of the securities they listed during the interval between nominations and the completion of the auction. The Treasury would select a subset of securities (e.g., 500) to be included in this particular auction, and it would decide on a total purchase budget. These decisions would be based in part on information obtained from the nomination process. For instance, the Treasury could decide not to purchase more than 50% of the face value of the assets nominated.

The Treasury planned to set "reference prices" reflecting estimates of the values of the various securities relative to a baseline security. A security with a reference price of (e.g.) 2 would therefore be deemed twice as valuable as the baseline security. Where possible, these reference prices would be calculated by combining the latest transactions data, other market information (e.g. Markit ABX indices), valuation models, expert opinions, and the outcomes of previous TARP auctions.

Each participating institution would be allowed to submit (a possibly limited number) of sealed bids.[10] Each bid consists of i) a unit price and ii) a number of units of a specific security. The Treasury would then normalize each bid by dividing the unit price bid by the reference price corresponding to that security. As a result, the bids for different securities are homogenized and can be compared. For instance, bids on a security with a reference price of 2 would be divided by 2 in order to make them comparable to bids on the baseline security. The normalization process would therefore lower the price bid for securities with higher reference prices, and raise the price bid for securities with lower reference prices.

The Treasury would accept the bids with the lowest *normalized* bids first, moving up the list until its budget for the auction is exhausted. Under a discriminatory payment format, the winning bidders (i.e. those whose securities are purchased by the Treasury) would be paid their own bids. Instead, the Treasury opted for a uniform price auction, whereby all winning bidders for a given security would receive the same payment per unit.[11] To

[9] This section borrows heavily from Armantier, Asker and Vickery (2008).

[10] Some multi-round auction formats were also widely discussed in the press at the time (Ausubel and Cramton, 2008). For example, a reverse "clock auction" could be set up so that bidders would indicate quantities they wish to sell at current clock prices, which would be reduced sequentially until requested sales no longer exceeded a government purchase budget. One potential problem with clock auctions was the time required to complete a single auction, given the need to run many auctions per day with up to 100 securities in each auction. Another issue was the high volatility in the stock market in the fall of 2008, which could end up being the focus of news coverage of any ongoing auction.

[11] Three arguments were mentioned to support the uniform price format: i) It may provide some protection against the winner's curse (as the bidders with the most extreme predictions would not be penalized as much since they will most likely receive a price exceeding their bids if their securities are purchased by the

determine this price, the Treasury would first calculate the "market clearing normalized price" (i.e. the lowest rejected normalized price bid). This *common* cutoff normalized price, multiplied by the corresponding reference price then determines the "market clearing price" per unit the winning bidders for a given security receive. For instance, imagine the highest accepted normalized bid is 5, and consider a security with a reference price of 2. The owners of that security who submitted a bid below 10 (therefore corresponding to a normalized bid below 5) would sell their securities to the Treasury for a market clearing price of 5*2=10 per share. Observe that the market clearing price paid for a given security is the same, regardless of the winners' bids, but the market clearing price paid for securities estimated to be less valuable (i.e. with a lower reference price) would be lower. For instance, the baseline security in the previous example (with a reference price of 1) would be purchased for a market clearing price of 5. The market clearing price for each security could then be announced in order to facilitate trading in secondary markets.

The Treasury team realized that the effectiveness of this design relied crucially on the ability to come up with reference prices that reflect the relative security values. Although estimating relative values of MBS is arguably easier than estimating absolute values, the Treasury was confident this could not be done with perfect accuracy. Consequently, concerns were raised that bidders might exploit mispriced reference prices. To illustrate this possibility, consider a situation where the Treasury "overprices" a given security, i.e. the reference price is set above the actual relative value of the security. For instance, a reference price for a given security could incorrectly be set at 4 while its relative value is in fact 2. The bidders who own this security then possess an advantage compared to the owners of correctly priced securities. Indeed, all else equal, the normalized bids will be lower for this security (as they are incorrectly scored against a higher reference price), which make them more likely to be accepted in the auction. As a result, bidders may find it advantageous to exploit overpriced reference prices by submitting higher bids on those securities, which could be costly to the Treasury.

Several measures were proposed to mitigate the problem of incorrect reference prices. One approach would be to set limits on how much of the total purchase budget would be allocated to a specific security (e.g. 50% of the security's total face value in the nomination process). In addition, for a given security, the Treasury could decide not to purchase more than a certain number of shares from a given bidder. An alternative approach considered was to keep reference prices secret at the time of the auction. As a result, bidders would not be in a position to know and exploit mispriced securities. Finally, one aspect of the Holt and Goeree proposal involved reference prices that would be endogenized, e.g., modified ex-post as a result of the bids submitted.[12]

Treasury), ii) it may encourage participation from smaller or less informed bidders (who know that making a mistake may not be too costly since they are unlikely to set the market clearing price), and iii) it would generate a single market clearing price for each security, which might help with price discovery. The main argument in favor of the uniform price auction, however, was familiarity, as most financial institutions were experienced bidders under this format through Treasury auctions or the recently implemented "Term Auction Facility."

[12] Additional features were considered. In particular, to facilitate banks' needs for target levels of liquidity. Charles Plott suggested that bidders be allowed to submit a small number of "conditional bids," which are

The objective of the experiment described in the next section is to test i) how a Reference Price auction performs compared to a Grand Auction when reference prices are set correctly, ii) how efficiency in a Reference Price auction is affected when noisy references prices are announced before the auction, and iii) how keeping noisy reference prices secret at the time bids are submitted improves the auction's performance.

III. Experimental Procedures

The experiment consists of 4 treatments. As explained below, these treatments essentially differ in two ways: i) the type of reference prices used to score the bids and ii) whether the reference prices are revealed before or after the auction. Six sessions were run for each of the four treatments, using undergraduate students. The experiments were conducted in the Veconlab at the University of Virginia and the CIRANO's Bell Laboratory for Experimental Economics in Montreal.[13, 14] A session, which lasted about an hour and fifteen minutes, involved a group of 6 bidders and 6 securities, labeled A through F. There were 10 units (shares) of each security, and each bidder was endowed with 5 units of one asset, 3 of a second, and 2 of a third. The ownership pattern was balanced in the sense that, for each security, 5 units were owned by one person (who therefore had some market power), 3 units were owned by another, and 2 units were owned by a third person. This ownership concentration was intended to provide more opportunities for strategic bidder behavior, and therefore more of a stress test of the auction designs. Given the complex nature of the environment and the relatively unfamiliar reverse (low bids win) auction format, we decided to include some context, without being too specific. For example, the items being purchased were referred to as "shares" of "securities" instead of "mortgage backed securities," and the purchaser was referred to as the "government" instead of the "Treasury." This terminology will be used in the discussion of experiment results.

There were 8 auctions per session. As explained in the previous section, we consider sealed bid, common value, uniform-price auctions with reference prices. Prior to each auction, each of the assets was assigned a common value that was randomly drawn from a uniform distribution on a range [20, 80]. For each security they owned, bidders were given a private, individual-specific "signal" that was independently drawn from a uniform distribution centered at the true asset value, with a range of plus or minus 10. Since the true values were between 20 and 80, the range of possible signals spanned the range from 10 to 90, and bidders were told to think of these numbers as "dollars per

---

packages (each consisting of a security and a price) accompanied with the condition "if one bid in the package is accepted, then all other bids in the package are cancelled."

[13] All auctions were run with the Veconlab Reverse Auction program, which is listed under the Auctions menu on http://veconlab.econ.virginia.edu/admin.php.

[14] We also ran two additional treatments (6 sessions each) in which reference prices were not ratios, but instead, were equal to the government (absolute) value estimates. The results are quite close to the results obtained for ratio reference prices discussed in Section III below, so we do not include them here. The treatments with relative reference prices were run at the University of Montreal, and the treatments with absolute reference prices and unitary reference prices were run at the University of Virginia.

hundred of par value.” In some of the treatments we also assumed that the government received noisy signals about each security (absolute) values as random numbers drawn from a uniform distribution with a range of plus or minus 20 from the true security value. In other words, the government receives lower quality information than the bidders, at is signals are twice as noisy. Note also that we used matching sets of random draws for each treatment (akin to a common random number technique), i.e. each treatment was done with the same series of six random number seed values (one per session). Therefore, all sets of random draws (the security values, the bidders’ and, when relevant, the government’s value signals) were the same across treatments for the 48 auctions in a treatment. As a result, the experimental outcomes obtained in each treatment may be compared directly.

The government purchase budget was set at 2,000 (experiment dollars), corresponding two-thirds of the expected value of the 60 security shares for sale at an auction (3,000). This ratio is relatively high compared to what the Treasury was actually considering for the TARP auctions (i.e. no more than 50% of the total face value of the securities included in the auction). This choice was motivated in part by the fact that we wanted our subjects to make money in the experiment. Indeed, we believed that the relatively high government budget provided a better opportunity to sell at high prices in the experiment. This feature of the design therefore also makes our experiment a stress test of the Reference Price auction.

Bidders could submit “supply functions,” i.e. they may enter a bid price for each of the security units they own. At the time subjects submit their bids the following information is common knowledge: i) the distributions of security endowments across bidders, ii) the distributions of the randomly determined security values, iii) the conditional distributions of the bidders’ signals, and iv) in the relevant treatments, the conditional distributions of government signals. For each accepted bid, the bidder would earn the difference between the price received from the government and the true value of the security. Otherwise, earnings would be zero. This “profit-based” procedure provided bidders with a strong incentive to consider bids carefully, instead of just collecting high earnings by retaining security endowments, a behavior pattern that was sometimes observed in pilot sessions. It follows that each auction is a zero-sum game, so government losses are equal to the bidders’ profits.[15] The cash payout rate was 0.05 (dollars per dollar of experiment earnings), and they were generally in the $13-$30 range, with an average of about $23.

The four treatments conducted are summarized in Table 1. Observe that the treatments essentially differ in two ways: how the reference prices are calculated and whether or not they are announced before the auction. We now describe in more details each treatment.

[15] Note that this may not have been the case in practice. Indeed, some have argued that the securities were worth less to the bidders than to the government because i) the government may have been a more patient investor, or ii) some bidders faced a need for immediate liquidity. The assumption of a zero-sum game may easily be relaxed (without affecting the bidders’ behavior) in our experiment by assuming that the value of the securities for the government is $k > 1$ times greater than the value of the securities for the bidders. Although we did not run experiments with liquidity targets, the experiment software also permits a setup in which liquidity obtained by bidders is multiplied by a constant ($k > 1$) if it exceeds a specified target liquidity level.

Table 1. Treatments and Associated Reference Price Structures

Treatment	Reference Price Estimates	Reference Price Information
Grand Auction (no reference prices)	None	Irrelevant
Announced Noisy Reference Prices	Noisy	Announced
Secret Noisy Reference Prices	Noisy	Secret
Announced Accurate Reference Prices	Accurate	Announced

Treatment 1: Grand Auction. Bids for all securities are collected, ranked, and the government purchases the securities with the lowest price bids (without any regard for the type of securities purchased) until its budget is exhausted. The price of the lowest rejected bid (i.e. with the lowest price bid among all rejected bids) determines the market clearing price. A bidder is paid the market clearing price for each unit of security accepted by the government. This implies that all securities are purchased by the government at the same price regardless of type. Notice that the Grand Auction corresponds to a special case of a Reference Price auction in which all reference prices are set to 1, and hence are irrelevant. In this case, reference prices provide no basis for adjusting bids to reflect relative security values, so the low bids should come from owners of the more "toxic" low-value securities. Because of the adverse selection problem, this simple combined auction was expected to yield low value for a given government purchase budget. This treatment may therefore be interpreted as a lower benchmark against which the treatments involving reference prices will be compared.

Treatment 2: Announced Accurate Reference Price Auction.
This treatment is the same as Treatment 1, except that the government first divides bids for a security by an accurate reference price announced to bidders before the auction. The reference prices in this treatment are perfectly accurate indicators of relative value, even if security values are not known by government officials in an absolute sense. More specifically, the reference price for a security is equal to the ratio of that security's value to the value of a baseline security. By construction, this baseline security "A," has a reference price of 1. As explained in the previous section, the government purchases the securities with the lowest normalized prices. The normalized price of the lowest rejected bid (i.e. with the lowest normalized price among all rejected bids) determines the normalized market clearing price, which is determined recursively so as to exhaust the auction budget. The actual price paid by the government for a security is then determined by multiplying the reference price ratio for that security by the common normalized market clearing price.

In this ideal case where relative security values are precisely known by the government, the reference prices should create a level playing field by homogenizing the bids for securities of different values, and should therefore promote competition across bidders and securities. In addition, bidders in this treatment do not have the opportunity to exploit the information they possess about the securities they own in order to extract a rent from the government. This scenario is therefore expected to provide an upper bound on purchase efficiency, at least for single-round auctions, as measured by actual value purchased to the auction budget.

Treatment 3: Announced Noisy Reference Price Auction. This treatment is the same as Treatment 2, except that the government now relies on noisy reference prices. More precisely, the reference prices are determined as ratios of the government noisy signal for each security to the government noisy signal for a base security. Note also that this treatment may be considered an extreme case of an auction with noisy reference prices for at least two reasons. First, the government's information about the *absolute value* of each security is poor. Indeed, its signals are twice as noisy as the bidders. Second, the government's information about the *relative values* of the securities is even poorer. Indeed, by taking the ratio of its noisy signals the government compounds its errors. As a result, the reference prices the government uses to decide which security to purchase may be highly inaccurate, or even close to meaningless in some cases. In our experiment for instance, the reference price of one of the security is incorrectly set one ninth of its actual relative value.

As explained in the previous section, bidders could take advantage of "overpriced" reference prices when noisy reference prices are announced before the auction. In such cases, a bidder could find it profitable to bid less aggressively (i.e. by submitting a higher price). This strategy may permit the bidder to increase the market clearing price without affecting his chances to be a winning bidder. Compared to the previous treatment, the opportunity for bidders' manipulation in response to mispriced reference prices may therefore lead the government to overpay to acquire the same value.

Treatment 4: Secret Noisy Reference Price Auction. This treatment is the same as Treatment 3 except that the government does not announce its noisy reference prices before the auction. The bidders must therefore form an expectation about the reference prices the government will use to score their bids. The reference prices, the allocation mechanism and the pricing rules, however, are exactly the same as in Treatment 3. Observe that there are actually two differences between these two treatments. The first is strategic, as the bidders in the treatment without announcements do not know the reference prices against which their bids will be scored. The second is informational, as the bidders who do see announced reference prices are able to observe additional (noisy) information about the relative values that can be inferred from government reference prices.

By not announcing reference prices before the auction, the government should be protected against strategic reactions to overpriced reference prices, and hence, the value of securities purchased for a given auction budget should be higher with secret (but noisy) reference prices, as compared with announced (but noisy) reference prices.

III. Laboratory Results

Summary auction outcome measures are organized by treatment in Table 2. As is apparent from the top row, the government nearly exhausts its $2,000 auction purchase budget in all treatments. Note also that the government expenditure is the lowest and the most variable in the Grand Auction. The second row in Table 2 gives the average actual values of the securities purchased by the Treasury at an auction. Observe that the government operates the auctions at a loss in our experiment since its expenditure exceeds the total value of the securities purchased for all four treatments. Since this is a zero-sum game, subjects earn positive profits in all treatments. This result is not surprising given the relatively high budget we set for the government in the experiment in order to guarantee our subjects positive earnings.

As would be expected, the average value of assets acquired is the lowest with the Grand Auction and it is the highest with accurate reference prices, as shown by the first two columns in row 2. Next consider the last two columns for the treatments with noisy reference prices. As expected, the value purchased by the Treasury when announcing the reference prices before the auction is intermediate compared to the two previous treatments. In contrast with our prediction, however, the average value purchased appears to be slightly lower when noisy reference prices are kept secret at the time of the auction.

These treatment differences may be better appreciated when considering our measure of purchase efficiency, i.e. the ratio of the value of assets purchased to government expenditure in row 3. We can see that the Grand Auction produced an efficiency ratio of 76.6%, meaning that for every dollar spent, the treasury purchases for only 76.6 cents in value. In contrast, despite the relatively high government budget, the accurate announced reference prices auction is almost fully efficient, with an efficiency ratio of nearly 94%. In other words, if the Treasury can acquire precise estimates of the securities relative values, then, as hypothesized, a Reference Price auction would homogenize the different securities and produce a nearly efficient outcome. The outcomes of the Grand Auction and the accurate announced reference prices auction are therefore consistent with our predictions as the two treatments constitute two opposite benchmarks.

The two treatments with noisy reference prices produce intermediate efficiency ratios. This implies in particular that, because of its homogenizing properties, a Reference Price auction is superior to the Grand Auction even when the government has little information about the relative values of the securities. Observe, however, that, in contrast with our prediction, keeping the noisy reference prices secret at the time of the auction does not improve the auction efficiency. Below, we show that this result is not due to the fact that the auction with secret noisy reference prices underperformed. Instead, we find evidence that the auction with announced noisy reference prices performed beyond expectation. To sum up, the Reference Price auctions clearly dominate the Grand Auction, even when reference prices are extremely noisy, but keeping reference prices secret before the auction does not reduce the government losses.

Table 2. Average Auction Outcomes by Treatment (with Standard Deviations)

	Treatment 1 Grand Auction: No Reference Prices	Treatment 2 Accurate Announced Reference Prices	Treatment 3 Noisy Announced Reference Prices	Treatment 4 Noisy Secret Reference Prices
(1) Government Expenditure	1,912.354 (103.584)	1,955.041 (48.004)	1,934.915 (38.757)	1,943.155 (52.770)
(2) Value of Securities Accepted	1,465.125 (207.140)	1,830.250 (177.549)	1,694.167 (208.496)	1,659.500 (142.554)
(3) Efficiency (i.e. Value Purchased to Government Expenditure Ratio)	0.766 (0.106)	0.936 (0.087)	0.876 (0.102)	0.854 (0.068)

Figure 1 shows that treatment differences in terms of efficiency also hold up on an auction-by-auction basis, averaged over each of the 8 auctions per session. (Please ignore the “counterfactual” line, which will be explained below.) As expected, in all eight auctions, the Grand Auction with no reference prices is the least efficient, while the auction with accurate announced reference prices is systematically the most efficient. Another interesting feature of the figure is that the efficiency measures are fairly flat, thereby showing no evidence of learning or strategic adjustments, except for a possible decline in performance across auctions in the Grand Auction, which will be evaluated in more detail below. In other words, the dominance of the Reference Price auction in terms of efficiency is immediate, and does not require the bidders to familiarize themselves to the slightly more complicated auction design.

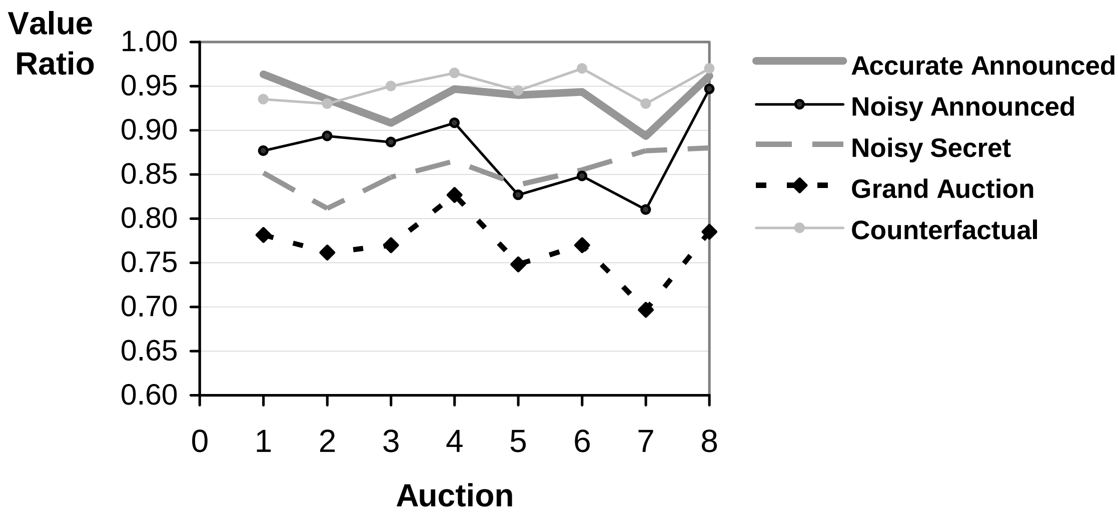

Figure 1. Efficiency by Auction (averaged over the 6 sessions)

The auction outcomes in the two noisy reference price treatments, and in particular the drop in efficiency compared to the benchmark auction with accurate reference prices, may be the result of both i) noise in the reference prices, and ii) strategic adjustments in the bidders' behavior to this noise. In an effort to disentangle these two effects we conduct a counterfactual exercise in which we take the observed bids that were submitted with noisy secret reference prices (Treatment 4) and score them against accurate (instead of noisy) reference prices, as was the case in the benchmark Treatment 2. The basic assumption underlying this counterfactual exercise is that subjects in the noisy unobserved reference prices treatment would have behaved in the same way had the *unobserved* reference prices been accurate instead of noisy. As indicated in Figure 1, the counterfactual exercise produces significantly higher efficiency measures in each auction than for the two noisy reference price treatments, and statistically indistinguishable from those observed when reference prices were accurate and announced (Treatment 2). This result suggests that the lower efficiency ratios obtained in the two noisy reference price treatments should be essentially attributed to noisy reference prices, not to adjustment of bidding behavior on the part of the subjects in the face of greater uncertainty about reference prices.

Since subjects in one session did not participate in another, we consider inferences based on averages for the 6 independent sessions in each treatment, which are shown in Table 3. Each treatment consisted of a session for each of the 6 random number "seed" values, so as to use parallel sets of random value and "signal" draws for each treatment. The session averages, arranged by seed, are shown in the second column of Table 3. For all 6 seed values, the efficiency ratios are lowest for the treatment with no reference prices (top row), and are highest for the treatment with accurate announced reference prices (bottom row). These differences are significant at conventional levels, using Wilcoxon signed-rank tests, as shown in the far right column of the table. Here, the Z and p values pertain to two-tailed tests between averages for a given treatment and averages for the treatment in the row just above it. Note, however, in 5 out of 6 cases, the efficiency ratios are *higher* with noisy announced reference prices than for noisy secret reference prices (p = .24 for a 2-tailed test).

Table 3. Efficiency Ratios per Session
(Averages are over all 8 auctions per session)

Treatment: Reference Prices	Session Efficiency Averages For Seeds 1-6	Treatment Average	Wilcoxon Signed-Rank Test*
None (unitary)	.79, .81, .79, .77, .72, .74	.766	
Announced, Noisy	.85, .86, .89, .92, .90, .83	.876	Z = –2.201 *P value* = 0.028
Secret, Noisy	.84, .84, .87, .83, .87, .87	.854	Z = 1.166 *P value* = 0.244
Announced, Accurate	.91, .91, .97, .89, .99, .95	.936	Z = –2.207 *P value* = 0.027

* Each test compares the distribution of the 6 ratios in the corresponding row with the distribution in the row above.

When purchase efficiency is measured in terms of the ratio of value purchased to government expenditure, the main results of this section can be summarized:

Purchase Efficiency Results: *In an exper.iment with uniform-price reverse auctions for heterogeneous securities, purchase efficiency is lowest when reference prices are not used (the Grand Auction) and it is close to full value when reference prices are accurate measures of relative security values. Intermediate efficiency measures are observed when the government only possesses noisy reference prices, but these efficiency measures are not improved by keeping reference prices secret to avoid strategic manipulation of bids for mispriced securities. The counterfactual exercise suggests that the lower efficiencies in the treatment with noisy reference prices are mostly due to the highly inaccurate reference prices, rather than to bidder responses to noisy reference prices.*

Next we consider the issue of price discovery, i.e. the extent to which the auction clearing prices for each security provide accurate information about their unobserved values. As shown in Figure 2, the median clearing prices are slightly above the underlying values for all four treatments. This bias may be considered relatively small given that the Treasury purchase budget was relatively high, at two-thirds of the expected values of all securities combined.[16]

[16] One interesting extension would be to determine the sensitivity of price discovery to the magnitude of the auction purchase budget. In other words, would the "demand shift" due to an even higher auction purchase budget cause more of an upward bias in clearing prices, or would competition at the margin of included securities keep the cutoff normalized price more or less stable.

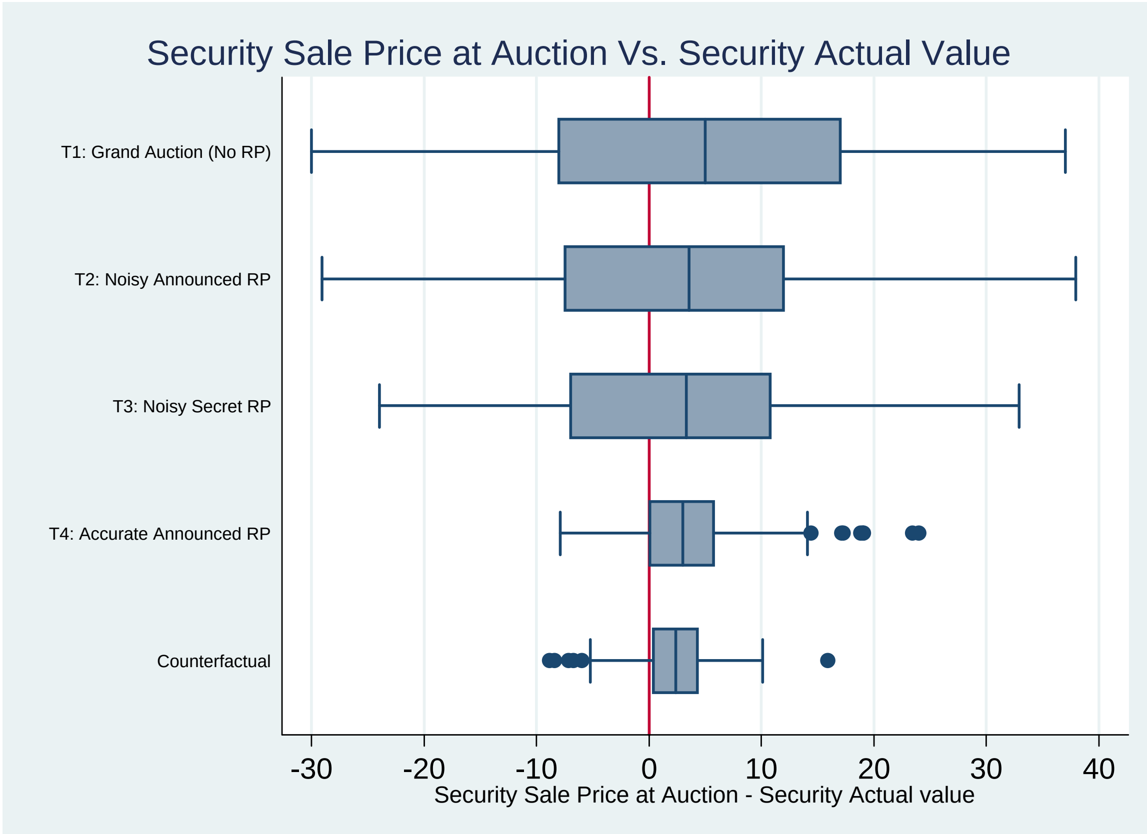

Figure 2. Security Sale Price at Auction vs. Security Actual Value

Note that, as expected, prices resulting from the Grand Auction are widely dispersed and essentially uninformative. In contrast, when reference prices are accurate (Treatment 2 and the counterfactual), the auction clearing prices are tightly distributed. There are no obvious differences between the two treatments with noisy reference prices; they are both better than the Grand Auction in terms of price discovery, but worse than the treatment with accurate announced reference prices and the counterfactual with accurate but secret reference prices. In other words, the extent to which security values are not accurately revealed in Treatment 4 (and by extension in Treatment 3) is due in large part to the fact that reference prices are measured with noise, not to bidders' behavior in these environments. These observations can be summarized:

Price Discovery Results: *Median auction clearing prices tend to be biased upward from underlying security values in all treatments, with the greatest dispersion in the Grand Auction (no reference prices) and the least dispersion in the treatment and counterfactual with accurate reference prices.*

The analysis of individual bidding patterns in the next section provides an explanation for the somewhat surprising result that purchase efficiencies and price discovery measures are not reduced with secret (but noisy) reference prices.

IV. Econometric Analysis

The analysis in this section is based on behavioral patterns that emerge from a series of reduced-form panel regressions of the form:

$$Y_{i,t,s,k} = \alpha X_{i,t,s,k} + \mu_i + \mu_t + U_{i,t,s,k},$$

where the subscript i = 1,..,36 indexes subjects in a specific treatment (6 sessions with 6 bidders each), t = 1,…,8 characterizes the auction number, s = 1,…,6 a security for sale at an auction, and k = 1,…,5 captures the order (from lowest to highest) of the bids submitted by a bidder for a unit of given security. So, k = 1 corresponds to the lowest bid submitted by a subject for a unit of a given security. Finally, μ_i is an individual random effect $(Var(\mu_i) = \sigma_i^2)$, and μ_t is an auction random effect $(Var(\mu_t) = \sigma_t^2)$.

Bidding Behavior

Table 4 shows the estimates for the reduced-form bid functions, where the dependent variable is the price bid by a subject to sell a unit of a given security. In all treatments, we find that, all else equal, subjects tend to submit higher bids when they receive a high signal (row 1) and for the last bids they submit for a given security (row 4). These results were expected. In particular, the second result simply reflects the fact that a subject’s bids for a given security are ranked in ascending order.

Table 4. Panel Data Model where the Endogenous Variable is the Price Bid

		Treatment 1	Treatment 2	Treatment 3	Treatment 4
		Grand Auction Unitary Reference Prices	Accurate Announced Reference Prices	Noisy Announced Reference Prices	Noisy Secret Reference Prices
1	Signal	0.901*** (0.016)	0.783*** (0.025)	0.709*** (0.023)	0.863*** (0.027)
2	Auction #	0.299*** (0.071)	0.447*** (0.121)	0.073 (0.154)	-0.227*** (0.074)
3	# of Contracts Owned	-0.441* (0.240)	-0.218 (0.421)	-0.775 (0.524)	-0.326 (0.247)
4	Bid #	3.452*** (0.542)	5.443*** (0.940)	5.568*** (1.136)	6.059*** (0.550)
5	(Bid #) * (# of Contracts Owned)	-0.309** (0.118)	-0.545** (0.202)	-0.580** (0.240)	-0.665*** (0.124)
6	Reference Price	—	6.647*** (0.454)	2.631*** (0.463)	-0.157 (0.216)
7	“Overpriced” Reference Price†	—	—	1.534* (0.885)	0.790 (0.608)
8	“Underpriced” Reference Price†	—	—	-5.894*** (1.785)	-0.162 (0.834)
9	Constant	-1.814 (1.455)	-6.847** (2.421)	3.676 (3.079)	1.200 (1.597)
10	Sigma *u*	5.548*** (0.550)	6.961*** (0.857)	7.897*** (1.333)	6.263*** (0.747)
11	Sigma *i*	7.543*** (0.093)	7.642*** (0.163)	10.103*** (0.197)	6.892*** (0.098)
12	Sigma *t*	4.125*** (0.103)	4.326*** (0.185)	5.921*** (0.187)	4.568*** (0.140)
13	*N*	2796	2876	2857	2872
14	Log Likelihood	-9283.1	-11216.7	-11674.1	-9690.3

† For security S, Overpriced (Underpriced) is set equal to $|A|*I_{A>0}$ ($|A|*I_{A<0}$), where A=Reference $Price_S$ – $Value_S$ / $Value_A$.
***, **, and *** indicate significance at respectively the 1%, 5% and 10% levels.
Standard errors are robust and clustered at the subject level.

We also identify different forms of learning across treatments. All else equal, subjects in the Grand Auction tend to bid less aggressively (i.e. submit higher bids) from one auction to the next (row 2 of Table 4). Likewise subjects in Treatment 2 (Accurate Announced Reference Prices) learn to increase their bids. Further analysis reveals that subjects in

Treatment 2 learn to submit higher bids for high-value securities (or equivalently, for securities with high reference prices). In contrast, subjects in Treatment 4 (Secret Noisy Reference Prices) lower their bids over time. They seem to learn that they must become more aggressive in order for their bids to be accepted.

In all treatments, subjects tend to send lower bids when they have some market power for that security (i.e. when they own more contracts of that security), although the effect is in general insignificant (row 3 of Table 4). Crossing the bid number and the number of contracts owned, yields negative and significant estimates, thereby indicating that bidders who own more contracts of a given security raise their prices from one unit of that security to the next at a slower pace (row 5 of Table 4).

Finally, we can evaluate the influence of reference prices on bidding for the three treatments with reference prices (row 6 of Table 4). When the reference prices are kept secret at the time of the auction (Treatment 4), we find no relationship between the bids submitted and the reference prices, as would be expected. In contrast, reference prices do have an effect when they are announced before the auction; bids are positively correlated with announced reference prices, as expected.

We are also in a position to evaluate how subjects adjust their bids in response to the errors the government makes when setting reference prices. To do so, we first define the Treasury mispricing as the absolute deviation between the Treasury's (noisy) reference price ratio and the true ratio of values (i.e. the ratio of the security's true value divided by the true value of security A). The variable "Overpriced Reference Price" (respectively "Underpriced Reference Price") is then set equal to this absolute deviation when the difference between the Treasury's reference price ratio and the true ratio of values is positive (negative), and zero otherwise. As expected, neither form of mispricing influences bids in Treatment 4 with secret reference prices, since the bidders are unable to observe the bias in advance. In that respect, keeping the reference prices secret does indeed provide the government some protection against incorrect reference prices. When noisy reference prices are announced (Treatment 3), however, we can see that, as expected, bidders take advantage of overpriced reference prices by submitting slightly higher bids (row 7). The effect however is only significant at the 5% level, and its magnitude is not significantly greater than in the noisy secret reference price treatment. Interestingly, our results also suggest that subjects bid substantially more aggressively (i.e. submit lower bids) when the reference price is underpriced (row 8). In other words, it appears that bidders who own a security that has been underpriced are likely to realize that they are at a disadvantage, and they may try to compensate by lowering their bids in order to remain competitive. Observing such behavior was somewhat surprising to us, as it not only reflects strategic sophistication on the part of our subjects, but also it was not anticipated when the Treasury team discussed the merits and potential drawbacks of the Reference Price auction. To summarize:

Bidding Responses to Noisy Reference Prices: *Bids are responsive to announced reference prices. Bidders also respond to perceived biases in government pre-announced reference prices. Namely, subjects increase their bids slightly when the reference prices* *are overpriced, and lower their bids substantially when the reference prices are underpriced. These two effects partially offset each other, which essentially explains why keeping the reference prices secret does not improve the auction purchase efficiency. In other words, one could say that the absence of significant difference between the last two treatments reflects the fact that the noisy announced reference price treatment performed above expectations, not that the noisy secret reference price treatment performed below expectations.*

Purchase Probabilities

Next we turn to selection issues, i.e. which securities tend to be purchased in the auctions. For the probit panel regressions in Table 5, the dependent variable is equal to 1 if a bid for a unit of a security is accepted, and 0 otherwise. For each treatment, a bid is more likely to be accepted i) when the bidder under-evaluated the security, i.e. the private value signal was below the actual security value (row 2), ii) when the bidder has market power, i.e. the bidder owns more units of a security (row 4), and iii) for the first bids submitted for a given security, as these correspond (by construction) to lower bids for a security (row 5). These results are consistent with intuition.

Likewise, as expected, lower value securities are significantly more likely to be purchased in the Grand Auction (row 1 of Table 5). Although, the magnitude is substantially lower, the same significant effect is found with secret noisy reference prices (Treatment 4). In contrast, the value of a security has no bearing on its probability of being purchased by the government for the treatments (2 and 3) in which reference prices are announced in advance. In other words, the use of reference prices appears to be effective in homogenizing securities of different values, but it is not fully effective when reference prices are only announced ex post.

In all treatments with reference prices (treatments 2, 3 and 4), the probability of a security being purchased is (all else equal) unrelated to its reference price (row 6). In both treatments with noisy reference prices, however, a bid is more likely to be accepted when a reference price is overpriced (row 7) and less likely to be accepted when it is underpriced (row 8). When the reference prices are kept secret (Treatment 4), this result only reflects the mechanics of the allocation process. Indeed, since behavior is the same regardless of the reference pricing errors, bids are automatically more likely to be accepted when a reference price is overpriced. When noisy reference prices are announced in advance, we saw from the estimated bid functions in Table 4 that bidders do indeed adjust their bids in reaction to biases in reference prices. The results obtained in Table 5 for Treatment 3 therefore reflect both a mechanical and a strategic effect. Observe in particular, that the magnitude of the effects of the two forms of mispricing is significantly lower when reference prices are announced in advance (Treatment 3) instead of being kept secret (Treatment 4).[17] Perhaps surprisingly, having noisy reference

[17] In fact, all else equal, incorrectly setting a security's reference price at 1.5 instead of 1 increases the security's chances of being purchased in the noisy announced reference price treatment by 9%. Conversely, incorrectly setting a security's reference price at 0.5 instead of 1 decreases the security's chances of being

prices has a greater influence on which securities are purchased when reference prices are secret.

Table 5. Panel Probit Model where the Endogenous Variable Equals 1 if Bid Is Accepted

		Treatment 1	Treatment 2	Treatment 3	Treatment 4
		Grand Auction Unitary Reference Prices	Accurate Announced Reference Prices	Noisy Announced Reference Prices	Noisy Secret Reference Prices
1	Security Value	-0.089*** (0.009)	-0.004 (0.004)	-0.006 (0.004)	-0.010*** (0.004)
2	Signal Bias (i.e. Signal – Value)	-0.082*** (0.014)	-0.104*** (0.010)	-0.058*** (0.006)	-0.066*** (0.010)
3	Auction #	0.005 (0.014)	-0.001 (0.019)	-0.024 (0.015)	0.031 (0.016)
4	# of Contracts Owned	0.242*** (0.055)	0.138*** (0.039)	0.137*** (0.037)	0.158*** (0.031)
5	Bid #	-0.312*** (0.035)	-0.266*** (0.042)	-0.194*** (0.021)	-0.298*** (0.029)
6	Reference Price	—	-0.013 (0.134)	-0.019 (0.021)	0.086 (0.107)
7	"Overpriced" Reference Price†	—	—	0.426* (0.204)	1.235*** (0.212)
8	"Underpriced" Reference Price†	—	—	-2.039*** (0.273)	-3.016*** (0.258)
9	Constant	0.536*** (0.071)	0.684** (0.277)	0.515 (0.171)	0.606* (0.212)
10	Sigma *u*	0.581*** (0.087)	0.569*** (0.071)	0.523*** (0.080)	0.436*** (0.057)
11	Sigma *i*	0.417*** (0.055)	0.512*** (0.050)	0.486*** (0.057)	0.461*** (0.039)
12	Sigma *t*	0.241*** (0.052)	0.193*** (0.032)	0.197*** (0.042)	0.275*** (0.049)
13	*N*	2880	2880	2880	2880
14	Log Likelihood	-881.2	-1492.3	-1562.1	-1401.2

† For security S, Overpriced (Underpriced) is set equal to $|A|*I_{A>0}$ ($|A|*I_{A<0}$), where A=Reference Price$_S$ – Value$_S$ / Value$_A$.
***, **, and *** indicate significance at respectively the 1%, 5% and 10% levels.
Standard errors are robust and clustered at the subject level.

purchased in the noisy announced reference price treatment by 32%. When reference prices are not announced, the analogous increase and decrease percentages are much larger, at 26% and 51%.

Purchase Probability Results:
In the Grand Auction the government tends to purchase securities of low value. When reference prices are announced, the purchase probabilities are largely uncorrelated with security values. When reference prices are noisy, announcing them before the auction helps soften (and not increase as hypothesized) the impact of mispricing on the values of securities that are purchased.

Government Losses

Finally, we consider the factors that affect the difference between the values of securities purchased by the government and the amounts paid. Recall that government losses correspond to bidder profits, and vice versa. The dependent variable for the panel regressions in Table 6 is the government losses on accepted bids. It is apparent from the reported estimates that, in all treatments, the government losses more money i) on low-value securities (row 1), ii) when the bidder overestimates the security's value (row 2), and iii) on the first bids submitted by a bidder for a given security (row 5). Observe in particular that, as expected, the extent of the first effect is significantly larger in the Grand Auction, without reference prices, where bidders are more likely to sell low-value securities. Note also that government losses in the Grand Auction tend to increase from one auction to the next. This result is a direct consequence of the increase in the submitted bids observed previously in Table 4.

Finally, consider the effects of using reference prices. When reference prices are announced (treatments 2 and 3), we find that all else equal, the government losses increase with the reference price (row 6), holding constant other factors such as the value of the security. This observation reflects the fact that subjects tend to bid higher in this situation. Finally, as one would expect, the government loses more when reference prices are overpriced (row 7) and more when reference prices are underpriced (row 8). Note that, although the magnitudes of both of these effects are larger in Treatment 3 with noisy announced reference prices than in Treatment 4 with noisy secret reference prices, the difference between the treatments is once again not significant. To sum up:

Government Loss Results: *In the Grand Auction, the government loses money on low value securities it purchases. This is much less the case in all three Reference Price auctions. The impact of incorrectly set reference prices on the government losses is not significantly affected by whether or not these noisy reference prices are announced or not.*

Table 6. Panel Data Model where the Endogenous Variable is the Government Loss for each Security Purchased

		Treatment 1	Treatment 2	Treatment 3	Treatment 4
		Grand Auction Unitary Reference Prices	Accurate Announced Reference Prices	Noisy Announced Reference Prices	Noisy Secret Reference Prices
1	Security Value	-0.872*** (0.016)	-0.118*** (0.014)	-0.185*** (0.033)	-0.152*** (0.027)
2	Signal Bias (i.e. Signal – Value)	0.081** (0.032)	0.122*** (0.034)	0.224*** (0.069)	0.245*** (0.059)
3	Auction #	0.239** (0.111)	-0.196* (0.128)	-0.327 (0.295)	-0.140 (0.182)
4	# of Contracts Owned	-0.120 (0.094)	-0.160** (0.069)	-0.004 (0.214)	-0.414 (0.265)
5	Bid #	0.124** (0.037)	0.061 (0.048)	0.199** (0.085)	0.731* (0.133)
6	Reference Price	—	6.672*** (0.450)	4.235*** (0.785)	0.599 (0.612)
7	"Overpriced" Reference Price†	—	—	4.710*** (1.560)	2.130*** (1.468)
8	"Underpriced" Reference Price†	—	—	-32.351*** (3.677)	-29.102*** (3.122)
9	Constant	35.546*** (1.430)	0.750 (1.032)	7.703*** (1.575)	10.735*** (1.638)
10	Sigma *u*	2.371*** (0.112)	0.966*** (0.141)	3.291*** (0.207)	1.947*** (0.284)
11	Sigma *i*	2.427*** (0.156)	2.332*** (0.205)	6.847*** (0.351)	6.959*** (0.364)
12	Sigma *t*	1.854*** (0.216)	1.629*** (0.174)	2.009*** (0.216)	1.998*** (0.235)
13	*N*	1745	1799	1666	1655
14	Log Likelihood	-4768.1	-4656.2	-5712.7	-5592.0

† For security S, Overpriced (Underpriced) is set equal to $|A|*I_{A>0}$ ($|A|*I_{A<0}$), where A=Reference $Price_S$ – $Value_S$ / $Value_A$.

***, **, and *** indicate significance at respectively the 1%, 5% and 10% levels.

Standard errors are robust and clustered at the subject level.

IV. Summary and Lessons for Structuring Future Auctions

The laboratory results indicate that running a combined auction for diverse and imprecisely valued securities, using “reference price” to normalize bids, would promote competition while providing some protection against adverse selection. As expected, in the Grand Auction without reference prices, the government purchases the worst quality securities and overpays. The efficiency level (i.e. the ratio of value purchased by the government divided by its expenditure) in that case is only 77% In contrast, when reference prices are perfect indicators of security values, the Reference Price auction works well in the sense that it enables the government to purchase securities almost at full value, with purchase efficiency of 94%. In a treatment where reference prices are noisy and pre-announced, efficiency drops to 88%. This result, however, may be considered remarkable as the government receives extremely noisy signals about the value of the assets in this treatment, and therefore sets nearly uninformative reference prices. Moreover, the pre-announcement of noisy reference prices does not reduce government purchase efficiency as compared with the case when these reference prices are kept secret. The absence of a substantial effect of announcing noisy reference prices may be explained by the combination of two factors that essentially offset each other. First, when announced reference prices are overpriced, subjects bid higher (less aggressively) leading to higher auction clearing prices. But when reference prices are underpriced, subjects bid lower (more aggressively) to stay competitive, leading to lower auction clearing prices. Although median auction clearing prices were biased upward from underlying security values in all treatments, the extent of this bias was several percentage points, although there was considerable dispersion in treatments with noisy reference prices. The lessons learned provide a framework for a more rapid response to a need for large-scale Reference Price auctions that may arise in a future crisis or in other contexts.

References

Armantier O. (1999) “Three Essays on the Empirical Application of Game Theoretic Models”, Phd Dissertation, University of Pittsburgh.

Armantier O., J-P. Florens, and J-F. Richard (1998) “Equilibre approximatif et regle intuitive: une application aux appels d'offres dans l'industrie spatiale,” *Economie et Prevision*, 132:179-190.

Armantier, O., and J. Vickery (2008) “An Auction Design to Purchase Toxic Asset,” internal memo, Federal Reserve Bank of New York.

Asker, J. and B. Cantillon (2008) “Properties of Scoring Auctions,” Discussion Paper, NYU.

Ausubel, L. and P. Cramton (2008) “Auction Design Critical for Rescue Plan,” *Economists’ Voice,* Vol. 5(5).

Cason, T., L. Gangadharan, and C. Duke (2003) “A Laboratory Study of Auctions for Reducing Non-Point Source Pollution.” *Journal of Environmental Economics and Management* 46: 446-471.

Cummings, Ronald, Charles Holt, and Susan Laury (2004) “Using Laboratory Experiments for Policy Making: An Example from the Georgia Irrigation Reduction Auction,” *Journal of Policy Analysis and Management,* Spring, 3(2), 341-363.

Goeree, J. and C. Holt (2008a) “Hierarchical Package Bidding: A Paper & Pencil Combinatorial Auction” *Games and Economic Behavior*, forthcoming.

Goeree, J., and C. Holt (2008b) “Standardized Uniform-Price Reverse Auctions A Multi-Market Auction Approach for the Treasury’s Rescue Plan,” draft proposal submitted to Treasury, October 9, 2008.

Holt, C., W. Shobe, D. Burtraw, K. Palmer, and J. Goeree (2007) *Auction Design for Selling CO2 Emissions Allowances Under the Regional Greenhouse Gas Initiative*, RGGI website**:** http://www.rggi.org/docs/rggi_auction_final.pdf

Holt, Charles A., William Shobe, and Angela M. Smith (2006) “An Experimental Basis for Public Policy Initiatives,” *Promoting the General Welfare: New Perspectives on Government Performance,* E. Patashnik and A. Gerber, eds., Washington, DC: Brookings Institution Press, 2006, pp. 174-199.

Stoneham, G., V. Chundhri, A. Ha, and L. Strappazzon (2002) “Victoria’s Bush Tender Trial: A Cost Sharing Approach to Biodiversity,” discussion paper, presented at the Inagural Sheep and Wool Industry Conference, Hamilton, Victoria, Australia.

Varian, Hal (2008) “How to Drink the Sub-Prime Kool-Aid,” *Economists’ Voice*, Vol. 5(5).

Appendix A. Instructions

Instructions, Page 1 of 6

- **Securities:** This is an auction in which you have the role of a bank with a portfolio of mortgage-backed assets, referred to as "securities." Owning different securities is like owning different stocks. Just as shares of stock in a given company are identical, units of a specific security are identical. These units will be referred to as "contracts."

- **Security Values:** There are 6 different securities, labeled A ... F. Each security has an underlying value that is the same for all bidders, but no bidder knows prior to the auction what this common value will turn out to be. You have the opportunity to sell some of your contracts to the government, but you will incur a loss if you sell contracts at prices that are below their values, which will be revealed after the auction.

- **Value Estimates:** For each security you own, you will receive a signal that is an imperfect estimate of the value of that security. Different bidders receive different signals about the value of the same security. Some bidders' signals may be above that security's true value, and some may be below. The average of all possible signal draws is equal to the underlying common value of the security (details to follow).

- **Purchase Budget for Combined Auction:** The government has budgeted a fixed amount, **$2000**, to purchase contracts for the 6 securities to be included in the auction. This budget will typically limit the number of contracts that can be purchased.

- **Reverse Auction:** This a "reverse auction" in which the government prefers to purchase from low bidders. You will be given the chance to submit an offer to sell your contracts for each of the different securities that you own. You will be bidding against **5 other bidders** in each auction.

- -

Instructions, Page 2 of 6

- **Reference Prices:** Since securities are different assets, some are intrinsically more valuable than others. Many of these securities are not actively traded, and the government has hired financial experts to estimate their values. Prior to the auction, each security has been assigned an official "reference price" that will be used to compare bids for different securities. The reference price is an estimate of the value of the security relative to the first security (A). For instance, if the reference price for security B is 2, then it means that the government believes that security B is twice as valuable as security A. And securities with reference prices below 1 are deemed to be less valuable then A. By construction, the reference price for security A is 1.

- **Comparisons:** The government is not necessarily interested in purchasing the securities with the lowest bid prices. Instead, the government is willing to pay more for securities that it believes to be more valuable. To compare bids across different securities, each bid submitted will be transformed into a **normalized bid** by dividing it by the reference price for that security.

- **Example:** For example, suppose that the reference price for a security is 2. Then a bid of $24 on this valuable security is the same as a bid of $12 on security A (which always has a reference price of 1), since both bids result in the same normalized price: $24/2 = $12/1 = $12.

- **Purchase Decision:** The government will accept bids with low normalized bid prices, moving to bids with higher and higher normalized bid prices until the fixed budget is exhausted or until all bids are accepted, whichever comes first.

- **Reserve Prices:** The government also reserves the right to reject offers that it deems to be unacceptably high, i.e. the maximum payment is capped at **$100** for each security.

- -

Instructions, Page 3 of 6

- **Cutoff Normalized Price:** As just explained, offers with lower normalized bid prices will be accepted, and offers with high normalized bid prices will be rejected. The cutoff normalized price is the lowest normalized bid price among the rejected bids. This cutoff is like a market-clearing price, and it determines the price paid for all of the bids that have been accepted. To summarize, bids with normalized bid prices below the cutoff are accepted, and those at or above the cutoff are rejected.

- **Uniform Sale Price:** For a given security, the government will purchase all accepted contracts at the same price: **Uniform Price = (cutoff normalized price)*(reference price)** for each contract of that security. Note that the price paid for two different securities will differ, since the reference prices of the two securities will generally be different.

- **Example 1:** Suppose that there is a bid price of 50 for a security with reference price of 2, and a bid price of 80 for a security with a reference price of 4. These bids have normalized prices of 50/2 = 25 and 80/4 = 20 respectively, so the 80 bid (with the lower normalized price) is more likely to be accepted. If this is the only accepted bid, the normalized price of 25 for the rejected bid of 50 is the cutoff normalized price. Thus contracts associated with the accepted bid of 80 would be sold for 25 times the reference price of 4 or for 100 per contract.

- **Earnings:** The amount you receive for a sale is at least as high as your bid, since all accepted bids have normalized prices that are at or below the cutoff normalized price. If a bid is not accepted, you earn nothing for those contracts. If a bid is accepted, the difference between the sale price and the value of the security will be added to you earnings. Thus you will be penalized if you sell below value, and you will earn more to the extent that your sale is above value (revealed after the auction).

- -

Instructions, Page 4 of 6

- **Example 2:** Suppose there are bids on 4 different securities, some with high prices and some low, but the normalized bids are ranked at 8, 9, 11, and 14. If the auction budget is such that only the bids with the two lowest normalized bid prices are accepted, then the cutoff normalized price for the first rejected bid is **11**, and each of the two bids with low ratios result in sales for amounts that equal 11 times the relevant reference prices. If the accepted bids were 16 with reference price 2 and 9 with reference price 1, then the bid of 16 will result in a sale at 22 = (11)*(2), and the bid of 9 will result in a sale at 11 = (11)*(1).

Bid Price	Reference Price	Normalized Bid	Sale Price
16	2	8 = 16/2	22 = (**11**)*(2)
9	1	9 = 9/1	11 = (**11**)*(1)
11	1	**11** = 11/1	(no sale at cutoff)
7	0.5	14 = 7/(.5)	(no sale)

- **Note:** What you receive for an accepted bid is not affected by the bid price. But bidding too high can be risky since a high bid is less likely to be accepted. Conversely, if a low bid is accepted, you will receive an amount at least this high, but bidding too low can be risky since the amount received could be as low as what you bid. You **incur a loss** if you sell contracts below their values (revealed after the auction).

- **Earnings Example:** If the true security values for the bids in the top two rows of the table turned out to be 20 and 10, then earnings would be $22 (sale price) - $20 (value) = $2 for the bid in the top row and $11 (sale price) - $10 (value) = $1 for the bid in the second row. Earnings are 0 for bids that are not accepted (bottom two rows). But if the values for the securities in the top two rows were 24 and 12, then earnings for these bids would have been negative: $22 - $24 = -$2 and $11 - $12 = -$1.

- -

Instructions, Page 5 of 6

- **Values:** The table below provides an example of your signal values and contract holdings for each of the 6 securities. The signal values are estimates and may turn out to differ from the actual values that are realized after the auction.

- **Value Range:** All security values are between $0 and $100, so you can think of bidding in terms of "dollars per hundred of par value" and the contracts as being par value amounts.

- **Value Distributions:** The value of any given security will be drawn randomly from a range between lower and upper bounds, with all values in that range being equally likely. For example, if the range is from 20 to 80, then the value of the security may be 20, 21, ... 80, as if the value were determined by a hard spin of a roulette wheel with stops labeled for each possible value in this range. The upper and lower bounds for the security value ranges will be shown in a table that you can look at while selecting your bids.

- **Signal Values:** Your signal for a given security is drawn randomly from a range between **$10.00** below the true value to **$10.00** above the true value. All signals in the interval are equally likely. The signals received by others about the value of a given security will be drawn randomly in the same manner. The range of possible signal draws for a more valuable security will be higher than the range for a less valuable security.

- **High or Low Signals:** If you have a high signal value for a given security, it could be because your signal draw was high in the interval of possible signals, or it could be because the signal draws for all bidders are high relative to those of other less valuable securities. Conversely, all bidders' signal draws will tend to be low for a less valuable security.

- **Reference prices:** For each security, the government receives a signal drawn from between **$20** below the true value and **$20** above the true value. The government then calculates the reference price by dividing its signal for each security by its signal for security A. For instance, if the government signals are 60 and 30 for securities A and B, then the reference prices are 1 (=60/60) for A and 0.5 (=30/60) for B. Note that a higher reference price will generally correspond to a more valuable security, but reference prices are subject to error since the government receives imperfect signals about security values.

Security Code	Reference Price (per contract)	Contracts Owned	Value Signal
A	$*.**	**	$*.**
B	$*.**	**	$*.**
C	$*.**	**	$*.**
D	$*.**	**	$*.**
E	$*.**	**	$*.**
F	$*.**	**	$*.**

- -

Summary Page

- **Bids** specify the price at which you are offering to sell contracts of a specific security.

- **Different bids** (offers to sell) may be submitted for different contracts of a given security. Each bid specifies the security and the bid price for a single contract.

- In order to compare bids for different securities in a combined auction, all bids are divided by pre-announced **reference prices** and the resulting normalized bids for all securities are ranked together; those with low normalized bids are more likely to be accepted.

- A **fixed budget** of **$2,000.00** is used to purchase the securities to be accepted, beginning with the lowest normalized bids and working up sequentially until the budget is exhausted (with ties at the cutoff decided at random).

- The **cutoff normalized bid** for the first rejected bid is used to determine the sale price for all bids with normalized bids below the cutoff. This cutoff normalized bid is the *same for all securities*.

- **Sale Price = (Cutoff Normalized Bid) * (Reference Price)**, so all successful bidders will receive the same amounts per contract of a given security, regardless of their actual bid amounts.

- The **sale price** for an accepted bid is at least as high as the bid, since the normalized bid for an accepted bid will be at or below the cutoff normalized bid.

- **Earnings** equal the difference between the revenues from contracts that are sold and their values (revealed after the auction). You earn nothing on contracts not sold.

- **Values** for each security will be drawn randomly from a range that is between $20 and $80.

- **Signals** for each security will be taken from a range that is within plus or minus $10.00 of its actual value.

- For each security, the government receives a signal that is within plus or minus $20.00 of its true value. Dividing this signal by the government's signal for security A gives the **reference price** of the security.

- There is only a **single round** of bidding in the auction. Bids may not be revised once submitted.

- There will be **8 auctions**, and your security values and signal value estimates will be randomly regenerated for each new auction. The program will keep track of your total earnings for all auctions.

- **Special Earnings Announcement:** Your cash earnings will be **5%** of your total earnings at the end of the experiment.