# An experimental investigation of repeated auctions and secondary market trading in emissions markets with bankable allowances

Matthew Olson[1,2], Karla Hoffman[2], Daniel Houser[3]

[1]The MITRE Corporation, McLean, VA 22102
[2]George Mason University, Department of Systems Engineering & Operations Research, Fairfax, VA 22030
[3]George Mason University, Department of Economics, Fairfax, VA 22030

We report data from laboratory emissions allowance markets in which allowances do not expire and can be banked between compliance periods. These periods consist of a sealed bid auction, trading, and then compliance. The markets consist of multiple sequential compliance periods. We observe (1) market prices reflect an expectation of future market prices, not underlying equilibrium; (2) allowance banking increases with uncertainty; and (3) the secondary market - not the auction - is the primary mechanism of overall market efficiency. Contrary to our hypotheses, we also find (4) no efficiency difference resulting from the use of a uniform price or discriminative price auction and (5) no price or efficiency differences resulting from differing bid reporting rules after the auctions.

Keywords: cap and trade, auction, double auction, price discovery, disclosure rules, experimental economics

# 1 INTRODUCTION

US federal responses to climate change may feature a pollution allowance (*cap and trade*) market as a primary means of reducing greenhouse gas emissions. The House passed the American Clean Energy and Security Act of 2009 (HR 2454) on June 26, 2009. Senate versions are under consideration.

We investigate the price discovery and efficiency effects of both the auction design and its associated reporting rules in the context of the ongoing secondary market. The proposed federal regulation draws primarily on the history of the federal sulfur dioxide ($SO_2$) market and the Regional Greenhouse Gas Initiative (RGGI) market. We consider aspects of each with our laboratory markets and show possible difficulties with these market structures. To close the paper, we suggest future investigation into an alternative market design that we hypothesize reduces practical regulatory burden and promotes market efficiency.

Our laboratory markets include sealed bid auctions interspersed with secondary market trading and compliance periods. In our markets, we use the auctions to disburse at least a quarter of the allowances into the market. We observe (1) market prices reflect an expectation of future market prices and are insensitive to banked allowance inventories that depress the competitive equilibrium; (2) allowance banking increases with uncertainty; and (3) the secondary market - not the auction - is the primary mechanism of overall market efficiency. Contrary to our original hypotheses, we also find (4) no efficiency difference resulting from the use of a uniform price or discriminative price auction and (5) no price or efficiency differences resulting from differing bid reporting rules.

# 2 BACKGROUND

We begin with a short review of the literature concerning the institutions comprising an allowance markets: continuous double auctions and single-seller discrete auctions.

## 2.1 Continuous double auctions

The continuous double auction and its price formation properties has been the focus of much work. The literature contains two threads of interest here: repeated, separate markets and multiple round durable goods markets. In the first case, endowments – but not subject earnings – are reset at the start of each trading period. The ending condition of one period does not affect the endowments at the start of the next. In the durable goods case, the tradable tokens pay a dividend at the end of the period and remain in the subjects' possession to start the next trading period.

Friedman (2010) provides a summary of laboratory markets in which goods are not durable. Most relevant here is the fact that in a stationary, repeated environment, even if the number of buyers and sellers is few, subjects quickly converge upon the equilibrium price defined by the induced values. Gjerstad (2007) notes this convergence occurs primarily during trading, not between periods, and, in fact, the price of the first trade in a period is more a function of the trade prices in the immediately previous period than it is of the induced values.

Allowance banking between compliance periods makes the durable goods studies relevant to this particular environment. Smith, et al. (1988) provide the seminal results here. They report the ease with which laboratory subjects establish price bubbles for durable assets when a common, random (drawn from a known distribution) dividend is paid to the holder of each tradable token at the end of each trading cycle. In these bubble instances, trading prices increase to a point well above the expected value of future dividend payment and then collapse before the market draws to a close. They cannot say why bubbles are so easy to establish in such environments, but note the tendency to form bubbles decreases with subject experience.

Our markets differ in that our traded assets, while infinitely bankable, have private values and must be spent to generate new wealth in the market. Once spent, the asset is removed from the market. Franciosi, et al. (1999) use a market and assets similar to ours. They, however, use the two-sided revenue neutral $SO_2$ auction design, and they conduct trading before the auction in each compliance period. In the markets that allow banking between compliance periods, they show strong correlation between auction and secondary market prices, while showing also some (but not conclusive) tendency towards bubble formation.

### 2.2 Single seller discrete auctions

Auction design within the context of a continuous secondary market also spans the auction literature. Compliance entities participating in the market will estimate a private value for an allowance, but the value is likely correlated with the private values of similar firms. The secondary market introduces a common value element as the market price represents a lower bound on the value of allowances for the market participants. Since only a portion of the period's allowances are auctioned at any one time (either because multiple auctions are held each compliance period or because a large fraction of the market is allocated for free), we must consider sequential models as well.

Kagel's (1995) survey compares theory and experimental results in affiliated private value and common value single shot auction. In particular, he notes the systematic overpayment in common value auctions in the laboratory. If a winner's curse is not present, greater information availability prior to the auction further increases bids. If a winner's curse is present, the added information induces a more realistic assumption of the common value, and bids are lower than otherwise.

Krishna (2002) discusses sequential private value auctions in his survey. With bidders who have a private value for exactly one unit each, the equilibrium bids for a series of one unit sales do not depend on the information revealed from previous auctions. Ortega-Reichert (1968), Hausch (1986), and Mezzetti, et al. (2008) show that under affiliated value models with either multi-unit demand or multi-unit sale, equilibrium strategies differ with the information reported from previous auctions.

The link between sequential auction models and the pollution markets is served by models that include both an auction and then a subsequent aftermarket. Haile (2003) notes the difference between these resale models and common or affiliated value models that do not allow resale is the fact that the common price element is endogenously determined in the resale environment. He also notes that incentives to signal in the original auction exist in such a market, but the magnitude and even direction of the deviation from true value depends on the details of the auction and the secondary market design as well as the information reporting procedures. Goeree (2003) also shows that the structure of the secondary market helps determine the direction of the signaling in the auction. Further, when bidders have the incentive to understate values, a separating equilibrium may not exist, and the auction becomes inefficient.

### 2.3 Current legislation and markets

The market described in the HR 2454 is based primarily on two markets: the US sulfur dioxide ($SO_2$) market and the Regional Greenhouse Gas Initiative's (RGGI) greenhouse gas market. The $SO_2$ market serves as the model for general market structure and administration. The quarterly uniform price auctions mirror the RGGI design. Though the bill specifies an auction design, it allows the use of an alternative design, "… if the Administrator determines that

[it] would be more effective, taking into account… costs of administration, transparency, fairness, and risks of collusion or manipulation," (HR 2454 §791(c), 2009).

### 2.3.1 Acid rain – $SO_2$

Compliance in the $SO_2$ market is annual with an auction at the start of the compliance year and an on-going secondary market. Allowances do not expire, and they are defined by a vintage which specifies the first year that it can be used for compliance. The vast majority of allowances (~97.2%) are given for free to the regulated community; the remaining 2.8% are auctioned. Half of the auctioned amount is sold seven years before it can be used for compliance. The remaining half is sold at the beginning of the compliance year in which it can first be used.

The $SO_2$ auction is a two-sided, revenue neutral, discriminative price, sealed bid auction. After each auction, all bid information is made public; bidders are identified by name, quantity, and price. Cason (1993) shows analytically that the seller mechanism introduces inaccurate price signals and inefficiency. Cason and Plott (1996) experimentally show a uniform price variant to improve price discovery and market efficiency. Brookshire and Burness (2001) show analytically that the revenue-neutral mechanism also generally sends incorrect price signals to the secondary market. The Chicago Board of Trade, having conducted the first three auctions for EPA, in 1996 recommended a switch to a two-sided uniform price design, arguing likely improvements to price discovery, revenue, and efficiency (EPA A-96-19, 1996).

In practice, the $SO_2$ auction functions as a one-sided auction since private sellers have preferred to exchange in the secondary market, making moot the problems with the seller mechanism. Additionally, within one year of trading (and prior to the first compliance deadline), bids in the auction were a reflection of secondary market prices (Joskow, et al., 1998), making moot any problems with the auction's price discovery characteristics.

### 2.3.2 RGGI

The Regional Greenhouse Gas Initiative (RGGI) was formed by ten Northeastern and Mid-Atlantic States in the absence of federal action on GHG emissions. Compliance periods are three years. The first regulates 2009-2011 emissions. Allowances do not expire. All ten states have agreed to auction at least 25% of their allowances; six have promised to auction 100%.

The allowance auctions are conducted quarterly with the first occurring in September 2008. RGGI governs primarily electricity generators, so the quarterly auctions allow firms to simultaneously consider the complementary electricity auctions. The frequency of the auction is intended to provide liquidity to the market and to reduce the size of required capital reserves without creating undue administrative burden.

The auction design is a single-sided, uniform price, sealed bid design with a minimum reserve price. Like the $SO_2$ market, two uncoupled auctions are conducted simultaneously. One is for allowances of the current compliance period's vintage; the other is for the next compliance period. Unlike the $SO_2$ market, bid information is kept private in the RGGI market. The list of qualified bidders who declare intent to participate is released along with the clearing price, some aggregate statistics about the bid set and the quantities awarded to each winning bidder. Not disclosed are the identities of winning bidders or bid prices.

The original design recommendation (Holt, et al. 2007) discusses experiments that guided the reasoning behind the RGGI design. In some treatments they allow banking and employ a secondary market, but, unlike our markets, the secondary market is represented by a single uniform price call auction. They make no mention market price reflecting expected future price, rather than underlying supply and demand.

RGGI's actual auction announcements do reveal some doubt as to the discovery capability of the auction, however. The minimum reserve prices have been constant from the first auction to the current one (RGGIb, 2009), and the fourth auction announcement specifically states, "the Participating States have determined that there are not sufficient, reliable market data available to establish a valid current market price," (RGGIa, 2009).

# 3 EXPERIMENT

Our laboratory market was intended to capture the dynamics of the presumed greenhouse gas market. Here we describe our market structures, their deviation from their natural counterpart, our laboratory methods, and the input data.

## 3.1 Experimental environments

The laboratory market designs are simplified variants of the EPA's $SO_2$ market and RGGI's $CO_2$ auction. They are framed, however, as production environments. *Input units* take the place of allowances, and instead of compliance, subjects manage *production* in which they convert input units into *output units*, which represent new wealth in the market.

Each year is modeled as a cycle of four institutions: information update, a sealed bid auction, a secondary trading market, and production. Figure 1 shows a single cycle along with the time allotted to each institution within the cycle.

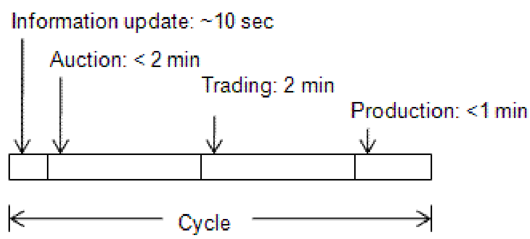


**Figure 1: Institution order and timing in a single cycle**

The cycle begins with the information update institution which does not require user interaction. Instead, all freely allocated allowances are distributed, and the subjects' assigned private values for output units are changed to introduce uncertainty into the experimental environment. These private values change at the beginning of each cycle, but remain constant through the cycle.

The auction follows the information update. Across all markets, the auction is a single-sided, sealed bid design that releases the year's remaining allowances into the market. Our treatments vary the payment rules, bid information disclosure rules, and the quantity of allowances sold in the auction each cycle.

The auction is followed by the secondary market that is modeled as a continuous double auction. Bids and asks are posted and transactions are executed one unit at a time. No bid or ask queue is maintained. Subjects can only submit offers that reduce the bid-ask spread, and once an offer has been beaten in the market, it is permanently removed from consideration.

Production is the final institution each cycle. Subjects are able to convert input units into output, the value of which is set by the values assigned at the beginning of a cycle. Each subject is limited to producing a maximum of five output units each cycle (though there is no limit on maximum inventory). One input unit is required to produce one output unit, and once converted into output, the input unit is removed from the market. Production is the only institution in which wealth is actually created. To model infinitely bankable allowances, input units do not expire; any unused inventory in a particular cycle is carried over into the next. However, input units that are not converted to output by the end of the market have no value to the subject.

### 3.2 Market simplifications

Our laboratory markets capture the salient features and incentives of natural greenhouse gas markets. At the same time, as in all laboratory investigations, they are unavoidably simplified. Our key simplifications are common to the literature in this area. We note our specific simplifications here. (1) In a natural market, emissions and allowance trading are concurrent. Only at compliance must the inventory of allowances be greater than the quantity emitted since the last compliance period; at any other time, a polluter may have already emitted more than it has allowances in inventory. In the laboratory market, this is not possible. The subjects first acquire input units and then use up to the number of input units in inventory to produce output. This simplification eliminates consideration of compliance penalties. (2) In our markets, all trading is spot trading and consists only of input units that can be converted to output as soon as they are released into the market. No derivative contracts exist. (3) All subjects in the laboratory market model polluters – all can use input units to generate wealth. In the natural market, non-compliance entities will participate in the market in an attempt to derive profit through exchange. In our market, profits can be made this way, but we have no subjects for whom it is the only way. (4) In a natural market, common information that is not provided by the market itself will be available. No external information is relevant to our laboratory markets.

### 3.3 Training procedures

Each subject participated in three experimental sessions (*Session 1*, *Session 2*, and *Session 3*), each of which lasted two to three hours. After the first two sessions, subjects were given a test to assess their understanding of the rules and the user interfaces. Subjects with adequate performance on the test were invited for a subsequent session. Session 3 markets are the primary source of data from which our conclusions are drawn.

Students earned a $5 show-up fee for all sessions. Session 2 and Session 3 markets featured performance payments; Session 1 markets did not. To ensure relatively constant average payouts across treatments, subjects participating in markets in which all input units were auctioned also received a flat fee to complete a session. This bonus was $5 for Session 2 markets and $10 for Session 3 markets. After both Session 2 and Session 3 markets, subjects were paid in cash at a rate one hundred experimental dollars to one real dollar. Session 2 earnings (not including show-up fee) averaged $26 per subject. Session 3 earnings (not including show-up fee) averaged $38 per subject.

We attempted to have subjects participate in only one treatment for their three sessions. This would allow them to learn the user interface and market rules concurrently with developing market strategies. For three of the five treatments this was true. For two treatments (Discriminative-Public-25 and Uniform-Private-25), some subjects participated in Session 1 markets that were different from their Session 2 and Session 3 markets. In all cases, subjects were provided with the same rules in Session 2 and Session 3 markets.

### 3.4 Laboratory procedures

To start a session, the monitor seated subjects randomly in the laboratory and then randomly assigned a username to each subject (Endowments had been assigned to the username as part of the experimental design and remained constant for the username across all laboratory markets in a given Session). Once given usernames, the subjects logged into the system, and the monitor distributed instructions. The monitor then read the instructions aloud. To improve subjects' comprehension, the reading was stopped after the explanation of each institution that requires user interaction (the auction, trading, and production). The subjects played the institution, and each was given a sheet of exercises. As subjects completed the exercises, the monitor spoke with each individually at their terminals to discuss any errors and to ensure understanding. The monitor did not collect the exercises; subjects could refer to them throughout the market and the test at the end of the session. After all subjects had been visited, the monitor did the exercises on the board for all to see. These explanations remained on the wall for the duration of the experiment. After the explanation to the group, the monitor resumed reading aloud the instructions for the next institution. The market proceeded without further interruption after the final bit of instructions.

### 3.5 Experimental design

Three two-level treatment variables are considered in the laboratory markets: sealed bid design (either uniform or discriminatory), bid disclosure rules (full disclosure or aggregate information), and the fraction of a cycle's allowances offered in the auction (25% or 100%).

Sealed bid design - The single-sided auction at the start of the cycle is either uniform or discriminative price. In both cases, revenue is returned to the seller, and no private sales are permitted during the auction phase.

Bid disclosure – After each auction, either all bid information (bidder identity and price) is reported or only the clearing bid (with identity removed) is reported. The full reporting case duplicates practices for the $SO_2$ market and the proposed federal greenhouse gas market.

Fraction of the allowances auctioned – Of the 24 allowances released into the market each cycle, either all are auctioned or only 25% are auctioned.

The three treatment variables are varied over five treatments. Three Session 2 markets and two Session 3 markets were run for each treatment. In all Session 3 markets and all but three Session 2 markets, eight subjects participated. In the remaining three Session 2 markets, only seven players participated, but in the cases where we invoke Session 2 data, these runs are excluded.

The uniform price auction serves as the auction design for the bulk of the markets. With the uniform price design, we have a two-replicate, 2x2 factorial design to assess the effects of bid disclosure policy and the fraction of allowances auctioned. We include the two markets featuring the discriminative price auction as they are more similar to the $SO_2$ auction, and they serve as a means of comparing our laboratory market to that natural market.

### 3.6 Input data

Subject endowments are constant across markets in a particular session. Session 2 markets last twelve cycles; Session 3 markets last sixteen. In both Sessions, twenty-four new input units are added to the market each cycle.

#### 3.6.1 Production values

At the start of a cycle, each of the eight subjects is given a new set of five descending marginal production values. Players are symmetric in that their value sets are drawn from the same distributions although the distributions vary between sessions and, in the case of Session 3 markets, between cycles.

In all sessions, maximum aggregate wealth results only when input units are converted to output in the same cycle in which they are first released into the market. In other words, banking reduces the overall wealth generated in the market. We did this to make a more confident a priori estimate of total subject payment. It had the secondary benefit of allowing the setting of a near-constant ceiling on the markets' equilibrium prices.

To assign the production values, we use three distributions each cycle. The first distribution sets equilibrium price (absent banking) by setting the 25$^{th}$ highest value (of the 40 created for the cycle). Since only 24 input units are released into the market each cycle and since banking is unprofitable, this $25^{th}$ value will never be exercised in a maximally efficient market. 24 values are then drawn from a distribution whose lower bound is strictly greater than the upper bound of the $25^{th}$ value. Finally, 15 more values are drawn from a distribution strictly below the realized $25^{th}$ value. These 40 values are then distributed randomly so that each market identity has 5 values per cycle.

The Session 2 production values are drawn from the same distributions for all cycles. The $25^{th}$ value is drawn from a discrete uniform distribution [e$37-e$40]. The high 24 values are drawn from a uniform distribution [e$41-e$100], and the remaining low 15 values are drawn from a uniform distribution [e$1 - $25^{th}$ value].

Session 3 markets include a values shift at the start of cycle 9 of 16 cycles. In all cycles, the high 24 values are drawn from [e$51-e$100]. In the first eight cycles, the $25^{th}$ value is drawn from [e$47-e$50], while in the second half of the market, the $25^{th}$ value is drawn from [e$27-e$30]. The result is a gap of over e$20 between the $24^{th}$ and $25^{th}$ values in the second half of Session 3 markets. This was intended to offer insight into speculative bidding with the bid disclosure policies and the relative price discovery capabilities of the two sealed bid options.

### 3.6.2 Free allocations

For the markets in which only 25% of the input units are auctioned, four of the eight subjects are given free input units each cycle. Five free input units are assigned to the owners of the two smallest sums of efficient production values. Four free input units are assigned to the owners of the third and fourth smallest sums of efficient production values. These subjects receive this quantity each cycle in the information update institution. We refer to them as *incumbents*; they model the entities that will receive free allowances under a grandfathering scheme. Subjects who do not receive free allowances are called *new entrants*. We concentrate the freely allocated input units in the hands of only four subjects to ensure each cycle still presents the opportunity for profitable trade after the auction.

## 4 PERFORMANCE MEASURES

Before proceeding into a detailed discussion of findings, we define terms for measuring market performance.

Equilibrium is the intersection of myopic supply and myopic demand. Since subjects have no explicit information about future value distributions, equilibrium is defined strictly in terms of current values and current inventories. Thus, as banked inventory holdings increase, the equilibrium price decreases.

A single player's demand is defined as the values that she cannot realize given her current inventory. As an example, Player A has three units in inventory. Since she has five descending marginal production values at all times and produces output in descending order of value, her demand is the bottom two elements of her set of production values. Similarly, player supply is the set of values that can be immediately satisfied from inventory holdings. Any inventory greater than five units is represented in the supply curve with value e$0 since subjects can produce no more than five output units per cycle. Input units offered in the auction also are considered part of the supply curve with value e$0.

In our laboratory markets, supply and demand rarely intersect at a particular value. In these cases, equilibrium is defined as the range between the last demand and supply value pair before curve intersection. If the curves are exhausted before they intersect, the considered range is between the last entry of the shorter curve and the entry with the same index in the other curve.

Banked inventory is the sum of input unit inventories after the production phase in a cycle. This equates to allowances retained by market participants after a compliance deadline.

Incumbent banked inventory ratio is the fraction of banked inventory held by incumbents over the total banked inventory at the end of a particular cycle.

Total market efficiency is the fraction of the total wealth generated in the market over maximum possible generated wealth. Since seller revenue is also included in the calculation of generated wealth, clearing prices above buyer values do not necessarily result in efficiency loss.

Holding efficiency ($eff_{hold}$) describes the wealth generating potential of the market's allocation at a particular time without considering production and banking. Given the sum of player inventories, $i_{current}$, the maximum holding value is the sum of the highest $i_{current}$ values at that time. The actual holding value is the sum of values that could be immediately realized if production were possible and all inventories were converted to output without further exchange. The holding efficiency is then the actual holding value over the maximum holding value.

Auction allocative efficiency ($eff_{auction}$) is a measure of the production values satisfied by the allocation resulting from an auction in which $q_a$ units are sold. The maximum possible auction value is the sum of the highest $q_a$ unsatisfied player values at the start of the auction. The actual auction value is the sum of the satisfied player values which corresponding to the $q_a$ highest bids. Auction allocative efficiency is then the actual auction value over the maximum value. We note that this measure does not consider the price bidders eventually pay. 100% auction allocative efficiency could still result in losses for bidders if they pay more than their value for the item.

Alternative auction allocative efficiency ($eff_{auction\text{-}alt}$) modifies the basic auction allocative efficiency value by including consideration of the equilibrium price. If the equilibrium quantity is strictly greater than $q_a$, then this measure says that any $q_a$-size set of the bids corresponding to the demand entries above the equilibrium price is fully efficient.

$$eff_{auction-alt} = \min\left\{1, \frac{\sum_{j=1}^{q_a} s_j}{\min\left\{\sum_{i=1}^{q_a} v_i \mid v_i \geq p_{eq}\right\}}\right\}$$

$s_j$ is a value satisfied in the auction. $v_i$ an underlying player value. $p_{eq}$ is the equilibrium price prior to the auction. The second term in the min expression is equivalent $eff_{auction}$ with the denominator replaced by the sum of the lowest $q_a$ underlying values that are greater than the equilibrium price. If the satisfied set consists exclusively of values greater than the smallest $q_a$ above equilibrium, this term will be greater than 1, thus the outermost minimizing operator.

## 5 HYPOTHESIS

We undertook this experiment with questions about the implications of the full bid disclosure, which is the policy for the $SO_2$ market and which we have assumed will be the policy under the presumed federal greenhouse gas market. We anticipated - consistent with resale models (Haile, 2003 and Goeree, 2003) - that subjects would use information disclosure rules to signal underlying values. Specifically, we hypothesized that the full disclosure of bids would encourage low bidding in the auctions as subjects sought to hide their true valuations from the other market participants. If this hypothesis were true, we would expect auction participants to pool bids at or below the underlying equilibrium price. The auction would, therefore, generate market inefficiency, and it would depress market price. Clearly, if such were the case, the government might carefully consider disclosure policy in future markets.

In the following section, we provide evidence that our hypothesis is incorrect. In particular, we find that subjects do attempt to communicate through the auction when bid curves are made public, but we find limited indication of an aggregate effect from bid privacy. In addition, we present surprising findings concerning price and risk response.

## 6 RESULTS

We now present our results in terms of five broad categories.

1. Price: Market prices reflect an expectation of future market prices and are insensitive to banked allowance inventories that may otherwise depress price.
2. Risk response: We find strong evidence that allowance banking is a response to uncertainty.
3. Primary auction versus secondary trading: The secondary market – not the auction – is the primary mechanism of overall market efficiency.
4. Auction design: With limited data, we find no difference in aggregate market behavior between the discriminative price auction and the uniform price auction.
5. Privacy: The full disclosure of bid information does not affect either overall market efficiency or seller revenues. It likely does not affect auction efficiency.

### 6.1 Prices

Our laboratory markets consistently yield both auction clearing prices and trade prices that are well above competitive equilibrium, reflecting instead an expectation of future market price. This separation of market price and the underlying equilibrium price is a consequence of banking, which itself is a response to uncertainty.

**Observation 1: In all treatments, auction clearing prices and trading prices reflect an expectation of future prices and are consistently above the underlying equilibrium range.**

Figures 2-11 are summaries of the Session 3 markets. They make plain the fact that market price is consistently higher than underlying equilibrium prices. They also argue against the price discovery properties of an auction once a market of this sort is established. Even in the three markets in which clearing prices and trading track equilibrium during the first eight cycles (Figures 3, 6, and 11), the auction at the start of the ninth cycle (when the underlying value of inefficient units falls) looks no different from the previous auctions. And after the shift, prices in both the auctions and trading phases fail to track equilibrium prices as closely as they had before the shift. The auction, if it provides discovery at the start of the market, does not provide that function in the middle of an ongoing market. Subjects bid and trade on expectation of future prices, not underlying value.

The finding is consistent with the literature. Our multiple cycles are stationary in the sense that the value distributions are constant over the first half of the market and then again over the second with a different distribution. The subjects do find a relatively constant market price as demonstrated in previous work with repeated, but separate double auction trading periods. With the banking mechanism, however, trading prices are above equilibrium, as is the case with previous durable goods experiments.

This result also strengthens the idea that price adjustment occurs during double auction trading rather than between trading periods – even if a period begins with a multiple unit auction, which considered by itself, theoretically should find the equilibrium price. Considering again Figures 2-11, at the start of cycle 9 when the gap is introduced in the value distribution, we see that the auction-clearing price in all but one market is higher than that of the subsequent double auction's average trade price. Whenever an adjustment (however minimal) occurs with the underlying value shift, we first see this adjustment in the double auction.

We also observe that the clearing prices are consistently above equilibrium prices. This result supports the existing common-value sealed-bid auction literature. That literature argues that one is likely to see a systematic overpayment in common value auctions either when the winner's curse is not a consideration or when bidders have greater access to information before the auction. Our environment satisfies both conditions. First, banking removes the winner's curse since allowance holders can hold them for later use or for resale if the immediate compliance period does not offer the opportunity for profitable use. And second, the secondary-market trading serves to increase information availability prior to subsequent auctions.

### 6.2 Risk response

**Observation 2**: **Allowance banking increases in response to increases in uncertainty about the value of inventory.**

Three pieces of evidence support this observation, but first, we recall the discussion of production values in section 3.6.1. In all three sessions, maximum efficiency can be achieved only if no units are banked. Value distributions are constant over Session 1 and 2 markets. For Session 3 markets, a values shock occurs at the start of the ninth cycles with cycles 1-8 using one distribution and cycles 9-16 using another.

<u>Session 2 vs. Session 3 markets</u>

At the outsets of Session 1 and Session 2 markets, subjects have no expectation of value distributions, changes to the value distributions, or market prices. In these Sessions, the distribution of production values remains constant over all cycles, and maximum efficiency can only be achieved if no units are banked. As such, these first two Sessions provide no opportunity for either speculative or hedging gain, and these subjects learn to spend more aggressively in Session 3 markets. The lesson is that as subjects become more comfortable in the environment (through the first two training Sessions) and as they begin to form expectations of future market conditions, they bank fewer allowances.

We show this by comparing the banked inventories at the end of each cycle across the Session 2 and Session 3 markets. In doing this, we consider only the first eight cycles. Session 2 markets conclude after twelve cycles, so by cycle 9, subjects begin shedding input units that have no value after the game. Session 3 markets include the demand shock at the start of the ninth cycle. Consideration of only the first eight cycles, therefore, allows comparison of strategies in constant value environments before any end of market effects can be felt.

Statistically, the null hypothesis holds that Session 2 banked inventories are less than or equal to those of Session 3. To test it, we apply a Wilcoxon rank sum test to each cycle, so that we have eight (correlated) tests. The first cycle yields a $p$-value of .0103; all other cycles have $p$-values less than .009. So even though the tests are strongly correlated, the weakest evidence is strongly persuasive. By creating an expectation of value stability over the whole course of Session 2 and the first half of Session 3, we remove perceived uncertainty, and banked inventories are smaller in Session 3.

<u>Response to value shift</u>

In Session 3, banked inventories are usually grown over the first four cycles and then trimmed through the eighth cycle as subjects liquidate inventories in response to an expectation of market stability. At the beginning of the ninth cycle, the demand schedule changes and uncertainty is introduced. As the market grapples with the shift, inventories are again grown. Of course, as the market again begins to settle and the close of the market looms, subjects begin to liquidate inventories. (Inventory not converted to output by the end of the market has no value.)

Statistically, we fit a linear model to compare the rate of inventory change for periods 5-8 and 9-12.

$$inventory = \beta_0 + \beta_1 c + \beta_2 x + \beta_3 cx + \sum_{i \in Runs-1} \gamma_i run_i$$

$c$ represents the cycle. To map the observations onto a common domain for regression, $c$ is simply the cycle number for cycles 5-8; for cycles, 9-12, it is the cycle number minus four. $x$ is a binary indicator noting whether the $c$ value has been altered. The interaction term shows the difference between the slopes (inventory change rate) of cycles 5-8 versus cycles 9-12. The summation term reduces overall variance by grouping the observations of each particular market.

The Session 3 markets yield a $p$-value of .002 for the null hypothesis stating the two banking trends are equivalent – strongly significant. The estimated slope for cycles 5-8 is indeed negative (-0.25 banked units/cycle) to show a consumption of banked inventory as subjects become accustomed to a constant value distribution. The estimated slope for cycles 9-12 is positive (0.91 banked units/cycle) to show the increased rate of banking after the value shift.

Inventory distribution (incumbents vs. new entrants)

In markets where only 25% of new input units are auctioned, it is the new entrant - not the incumbents - who are responsible for the banking. Statistically, we consider the incumbent banked inventory fraction at the end of each cycle. A signed rank test for each cycle tests the null hypothesis proposing that the median value of fraction held by incumbents is greater than or equal to 0.5. For the first cycle, the test's $p$-value is significant at 0.039; for all subsequent cycles, the $p$-value is at least half as small. As above, the tests are highly correlated, but even the weakest evidence for rejection is strong. Thus, we conclude that new entrants hold a larger fraction of the total banked inventory.

The disparity in banking rates can be traced to a difference in the uncertainty facing incumbent and new entrant subjects. Incumbents automatically receive enough input units to cover their maximum production capacity (or just one input unit shy of maximum capacity), so banking in consideration of their own values is unprofitable. Incumbents do, however, have the opportunity for speculative gain if new entrant production values and, therefore, market prices increase. New entrants share the same possible speculative gain if their own production values rise, but, unlike the incumbents, they also bear the risk of incumbent production values and, therefore, market price increasing. Thus, under the hypothesis that uncertainty increases banking, we should see new entrant firms banking more than incumbent firms. We indeed do observe this phenomenon.

### 6.3 Comparing auctions to trading

**Observation 3**: **The secondary market, not the auction, is the primary mechanism of overall market efficiency.**

As expected, when only 25% of a cycle's allowances are auctioned, the secondary market improves the wealth generating potential of the market (holding efficiency) over that achieved with just the auction. At the end of the auction, the incumbent subjects have not yet had the opportunity to sell their less profitable input units. Trading corrects this imbalance. In Session 3, of the 96 cycles (6 markets of 16 cycles each) only one cycle yielded holding efficiency after the auction that was greater than holding efficiency after trading.

The trend is weaker with the markets in which all input units are auctioned, but still holds. In 44 of 64 cycles (4 markets of 16 cycles each), the secondary market improves the holding efficiency resulting from the auction. Since this difference between the two institutions is only marginally correlated with the difference from the previous cycle (Durbin-Watson D = 2.1, Pr < D = .67), we model the observations as independent. Both a t-test and the signed rank test indicate that holding efficiency is greater after trading than after the auction (one-sided p-values less than .0001 for both tests).

Combined with Observation 1 in which we show the auction not to be a price discovery mechanism once the market is underway, the secondary market's consistent improvement of holding efficiency tells us that the secondary market is the primary mechanism of market efficiency and price determination. The auction may have effect at the market's start; in time, however, it does not. This finding is consistent with Joskow's (1996) analysis of the $SO_2$ market.

### 6.4 Auction design

Our experimental design is focused on the uniform price design as it is the most likely design for a future market, and so we compare the discriminative price and uniform price auctions only in markets in which 25% of the allowances are auctioned.

**Observation 4**: **The choice between uniform price and discriminative price auction design does not affect aggregate market performance when 25% of the allowances are auctioned.**

We compare the uniform and discriminative price auctions with three measures: overall market efficiency, seller revenue, and auction efficiency. We detect no difference for any of these measures – a result that fits squarely with Observations 1 and 3. The auction reflects the secondary market, and the secondary market is the primary means of market efficiency. The auction is superfluous, and thus its specific design does affect the aggregate market.

Market efficiency and revenue are cumulative effects, so we only compare across cycles, and in fact, only the final cycle should be considered to eliminate any strategy differences as subjects approached the end of the market. For total market efficiency, we consider all six Session 3 markets where only 25% of the input units are auctioned. At the end of the final cycle, a rank sum test yields a two-side p-value of 1; a general linear model yielded a p-value of .75 for the null hypothesis claiming no treatment difference. No prior cycle yields a non-parametric p-value less than .50 or a parametric p-value of less than .40. For total revenue, the difference is even weaker. At the last cycle, both tests yield p-values of 1, and no cycle yields a p-value less than .80 for either test.

The auction's ability to allocate input units to values above the equilibrium point is not affected by the choice of a uniform or discriminative price auction. We employ a general linear model similar to that in Observation 7 to identify any effect on alternative efficiency due to auction design. $R^2$ are extremely poor (0.02 for Session 3. 0.17 for Session 2), and neither test yields significance better than roughly 0.5.

### 6.5 Privacy

As mentioned above, we began this effort with questions as to the effects of the policy for bid disclosure after the auctions. Our findings, however, did not match our expectations. Here we present detailed results for our single positive finding related to privacy effects (Observations 5). We also present two negative findings (Observations 6 and 7), but limit detailed discussion to the (longer) working version of this paper (self reference omitted for peer review).

**Observation 5**: **The complete disclosure of bid information after the auction allows individual attempts at market manipulation.**

The markets present no macroscopic indications of market manipulation or collusion; however, questionnaires distributed after Session 3 markets indicate some attempts to use the auctions as a means of communication. For the markets in which subjects' auction bids were made public after the auction, subjects were asked

> *"When bidding in the sealed bid auction, did you consider the fact that your bids were public information immediately afterwards? If so, how did you alter your bidding strategy?"*

Two of the 56 subjects who participated in a final session that featured full bid disclosure explicitly indicated in their remarks that they attempted to communicate through the disclosure rules. The first one attempted to signal in the market featuring the uniform price auction and in which all input units are auctioned. The subject notes an effort to reduce the clearing price of the auction.

> *"I started bidding extravagantly on one or two units (to guarantee a buy @ clearing price) but I noticed that this made other players do the same, so I toned it down to lower the clearing price."*

By the fourth cycle in this market, the auction clearing price drops from initial highs and remains relatively constant until the end of the market. The subject is not able to coordinate any collusion to further suppress price, so he is likely mistaking cause for correlation as other players learn to bid less aggressively in the auction, but this does demonstrate intent.

The second - and stronger - example of communication via auction disclosure is from a market where 25% of the input units are sold with a discriminative price auction.

> *"Yes, by bidding in random cycles. Also by altering the buy and sell prices i.e. sometimes buy a unit for higher price to set a benchmark and then sell a few quick units (and it worked!)"*

Here, rather than seeking coordination, the subject pays to announce misinformation. If only the clearing bid is announced, this method of manipulation cannot be attempted. The purposely-high bids will never be conveyed to the market, and the subject will simply reduce his earnings without a chance for gain.

Seven other subjects (of the remaining 54) returned ambiguous free-form responses that may indicate some attempt either to hide private values by abstaining from the auction or to manipulate the auction price (both up and down).

These are anecdotal and self-reported evidence, and, again, we detect no macroscopic effects. Indeed, the multiple strategies revealed in the questionnaire responses seek to move auction prices both up and down, so, on the aggregate, they may tend to cancel. However, these responses are evidence of attempts to manipulate the market through the disclosure rules, and the government will have to balance the need for perceived transparency against the fact that the reporting will induce attempts at collusion and price manipulation.

**Observation 6: Bid privacy does not affect total market efficiency or seller revenue.**

Paralleling Observations 1, 3, and 4, bid privacy affects neither total market efficiency nor total seller revenue. The auction reflects the secondary market. The auction's details do not affect the overall market.

Since total market efficiency and total seller revenue are cumulative effects, we compare the values across markets at the end of each cycle. Further, we only consider the markets using the uniform price auctions for this observation, since we made no runs with discriminatory auctions and bid privacy. The result is that for both market efficiency and total revenue, eight observations are taken and considered in sixteen (highly correlated) statistical tests.

First we consider total seller revenue. An ANOVA (considering also the fraction of input units auction each cycle) across the eight markets shows that no cycle yields a p-value less than 0.25 for the bid disclosure treatment. Thus, we do not reject the null hypothesis which assumes no effect on total market efficiency due to auction disclosure rules.

The experimental markets give even less evidence to reject the null hypothesis that assumes no difference to revenue due to bid disclosure. Again applying a standard ANOVA across the eight markets for each cycle, the strongest evidence for rejection is a single cycle with a p-value greater than 0.45.

**Observation 7**: **Bid privacy likely does not affect the auction's allocative efficiency.**

We differentiate between market efficiency and auction's allocative efficiency in this observation. While market efficiency is a measure of generated wealth, allocative efficiency is a measure of the auction's ability to allocate input units to those valuing them most highly at the time of auction. A low allocative efficiency would indicate the auction provides an arbitrage opportunity with the secondary market as low-value bidders win units in the auction for later resale. A high allocative efficiency indicates high-value bidders actually submit the high bids and win units in the auction.

With this observation, we note that privacy rules do not, on the aggregate, affect the subject's ability to find arbitrage opportunities between the auction and the secondary market. Observation 5 shows clear individual attempts at arbitrage between the auction and secondary market, which if all attempted in the same price direction could conceivably affect efficient auction distribution, but we find no aggregate effect.

We analyze auction efficiency (both the standard and alternative versions) with the following regression model. The $y$ terms are indicators of the treatment. $y_{shock}$ is used only for Session 3 markets and notes whether the auction occurs before or after the value shift at the start of the ninth cycle.

$$eff_{auct,n} = \beta_0 + \beta_1 eff_{auct,n-1} + \beta_2 y_{privacy} + \beta_3 y_{fraction} + \beta_4 y_{privacy\times fraction} + \beta_5 y_{shock}$$

Table 1 shows results for both Session 2 and Session 3 markets. As discussed above, Session 2 markets were intended for training, but here they serve as a check against evidence from Session 3 that is statistically strong, but causes heartburn nonetheless.

| Auction efficiency measure | Session | Previous (alternative) auction efficiency | isPrivate | is100 | isPrivate × is100 | isAfterShock | $R^2$ |
|---|---|---|---|---|---|---|---|
| Standard | Session 3 | 0.22 (0.02) | -0.08 (<0.0001) | 0.09 (<0.0001) | 0.11 (<0.0001) | 0.01 (0.41) | 0.70 |
| Alternative | Session 3 | 0.20 (0.02) | -0.05 (0.0037) | 0.02 (0.1823) | 0.07 (<0.0001) | 0.02 (0.17) | 0.38 |
| Standard | Session 2 | 0.41 (<0.0001) | 0.01 (0.61) | 0.10 (<0.0001) | 0.03 (0.16) | | 0.46 |
| Alternative | Session 2 | 0.31 (0.0005) | 0.03 (0.26) | 0.01 (0.58) | 0.01 (0.77) | | 0.12 |

**Table 1: Auction efficiency – privacy and fraction auctioned - The top number in each cell is the estimated coefficient. The lower is the two-tailed p-value.**

Session 3 markets show that the efficiency of the previous auction, the bid disclosure rules, the fraction of input units auctioned, and the interaction of privacy and fraction are all strongly significant. Closer inspection of the data, however, reveals that the second replicate of the Uniform Private 25% markets is an outlier, in that clearing and trading prices are well above even the zero banking prediction of market price in the second replicate. Removing this outlier from the regression input, however, we again find the same variables significant.

Still skeptical, however, we apply the same model to the Session 2 runs for confirmation but consider only cycles 3-12 to avoid any difficulty associated with the learning process. Here, only the previous efficiency and the fraction of input units auctioned are significant.

We also note the low $R^2$ values for the models of alternative auction efficiency. This measure is more useful to the market designer as it describes the market's ability to allocate input units to any subset of true values above equilibrium. The $R^2$ values show the presented model, which is a function of the treatment variables, to be a poor description of alternative efficiency. Combined with the conflicting results of the two sessions and the outlier data in Session 3, we suspect that the null hypothesis of no privacy effect on auction efficiency is likely correct.

For practical purposes, even if bid disclosure policy does affect efficiency, we see in Observation 6 that it does not affect market efficiency, and we see in Observation 3 that the secondary market trading is generally the driver for market efficiency.

**Observation 8: The interaction of bid privacy and all input units being auctioned reduces trade volume in the secondary market.**

To assess effects on trade volume, we again turned to a general linear model, which includes the holding efficiency of the auction immediately prior to the observed secondary market instance, the trade volume of the previous secondary market, the treatments, and the value shock indicator.

$$vol_n = \beta_0 + \beta_1 eff_{hold-auct} + \beta_2 vol_{n-1} + \beta_3 y_{privacy} + \beta_4 y_{fraction} + \beta_5 y_{privacy \times fraction} + \beta_6 y_{shock}$$

Table 2 shows the results for both Session 2 and Session 3 markets. Session 3 markets show strong significance for the interaction term and the privacy term – both decreasing trade volume. Given our difficulties with the same variables and their effects on auction efficiency, Session 2 is used as a check. In these markets, the interaction term remains strongly significant; privacy by itself does not. We, therefore, are wary of claiming privacy alone affects volume, but we do conclude that as the fraction of the market auctioned increases, privacy provides stronger downward pressure on trade volume.

| Session | Intercept | Holding efficiency of prior auction | Trade volume of previous cycle | isPrivate | is100 | isPrivate × is100 | isAfterShock | $R^2$ |
|---|---|---|---|---|---|---|---|---|
| Session 3 | 20.05 (<0.0001) | -16.53 (<0.0001) | 0.18 (0.02) | -0.99 (<0.0001) | -1.00 (0.0002) | -1.28 (<0.0001) | 0.57 (0.01) | 0.62 |
| Session 2 | 6.82814 0.03 | -2.50215 (0.47) | 0.42630 (<0.0001) | -0.25593 (0.44) | -1.25642 (0.003) | -0.99263 (0.01) | | 0.50 |

**Table 2: Trade volume - The top number in each cell is the estimated coefficient. The lower is the two-tailed p-value.**

## 7 CONJECTURE, IMPLICATIONS, AND FURTHER WORK

The tendency of the participants in these mechanisms to hoard allowances and thereby keep prices higher than the equilibrium prices is the most practically distressing observation from the experiment. The cap-and-trade market is intended to aggregate the private abatement costs of the individual market participants to find the minimum total cost solution of satisfying the pollution constraint. In the process, market price should represent the regulated community's cost of abating the next unit of pollution. Our experiment suggests an inability of subjects in a market featuring periodic large auctions, continuous secondary market trading, and banking to find the market's marginal price. Rather than reflecting true underlying values, market prices in our environment seem to be more a response to price expectation and uncertainty. Despite consistently high prices, however, the relatively high holding efficiencies at the end of trading indicate that subjects are still able to use the high prices to determine the relative strength of their values and reallocate accordingly. This is indicative of a relatively healthy market. Government revenues are indeed higher than would be expected from underlying equilibrium prices – a finding that will be greeted with varying degrees of enthusiasm depending on one's relationship with the market, but the market itself is not destroying wealth. We caveat this assessment, however, with the facts that we have considered neither asymmetries between the market participants nor common information provided from sources external to the market.

Our laboratory markets, coupled with the observed results of the $SO_2$ and RGGI markets, call into question the basic structure of auctions in an allowance market since a primary objective of the auctions is price discovery. The original $SO_2$ auction may have helped price discovery early in the market, but came to mirror secondary market trading by the third auction. RGGI is still unsure as to what its market price should be. In our experiment, we see that once the market is established, the auction does not help identify the divergence of market prices from underlying use value, but does invite attempts at price manipulation. We also see that secondary market trading always improves the holding efficiency of a market after an auction.

We therefore conjecture that the emissions allowance market may not require periodic auctions. Instead, the government could release a steady stream of allowances into the secondary market. Economically, this may improve market efficiency by offering a guarantee of at least limited liquidity from the government even in the face of external shock. This may eliminate a degree of risk to the market participants, thus reducing the impulse for allowance banking, and in doing so, such a sales mechanism may deliver a market price that better represents the marginal cost of abatement. In terms of regulatory burden, the steady release of allowances obviates the government's cost of conducting an auction at all, and government revenues correspond exactly to current market prices. Such a market constitutes the next phase of our research.

## 8 CONCLUSION

We have presented the results of laboratory markets that model pollution allowance markets. The markets use bankable allowances in repeated cycles of a sealed bid auction, a secondary market, and compliance. We observe (1) market prices reflect an expectation of future market prices and are insensitive to banked allowance inventories that depress the competitive equilibrium; (2) allowance banking increases with uncertainty; and (3) the secondary market - not the auction - is the primary mechanism of overall market efficiency. Contrary to our original hypotheses, we also find (4) no efficiency difference resulting from the use of a uniform price or discriminative price auction and (5) no price or efficiency differences resulting from differing bid reporting rules.

Our findings call into question the basic structure of the proposed federal greenhouse gas market in which an active secondary market is expected and periodic sealed bid auctions are used to disburse some fraction of the allowances into the market. The auctions are not the primary driver of efficiency in such a market, and they fail to find prices that represent underlying use values and banked inventories.

In response, we have proposed and are currently experimenting with an alternative market design in which the large periodic auctions are replaced with a steady stream of allowances released directly into the secondary market. Economically, we hypothesize that the steady release of allowances will improve liquidity, thereby lowering the market participants' perceived risk and, therefore, lowering banked inventories. In turn, this could lead to both improved market efficiency and deliver lower market prices that more accurately represent abatement costs. Administratively and politically, the mechanism may be superior since the cost of conducting the auction is eliminated and since government revenues will more closely reflect market conditions.

## 9 REFERENCES


Brookshire, David S., and H. Stuart Burness. "The Informational Role of the EPA $SO_2$ Permit Auction." *Journal of Regulatory Economics* 20.1 (2001): 43-60.

Cason, Timothy N. "Seller Incentive Properties of EPA's Emission Trading Auction." *Journal of J. Environ. Econom. Management* 25 (1993): 177-195.

Cason, Timothy N., and Charles R. Plott. "EPA's New Emissions Trading Mechanism: A Laboratory Evaluation." *J. Environ. Econom. Management* 30 (1996): 133-160.

Franciosi, Robert, R. Mark Isaac, and Stanley S. Reynolds. "Experimental Research on the EPA's "Two-Tier" System for Marketable Emissions Permits." *Emissions Permit Experiments*. Ed. R. Mark Isaac & Charles Holt. Stamford, CT: JAI Press, Inc., 1999: 163-180.

Friedman, Daniel. "Preferences, beliefs, and equilibrium: What have experiments taught us?" *Journal of Economic Behavior & Organization* 73 (2010): 29-33.

Gjerstad, Steven. Price Dynamics in an Exchange Economy. Purdue University Economics Working Paper No. 1205, Purdue University, Department of Economics. November 2007. <http://www.krannert.purdue.edu/programs/phd/Working-paper-series/Year-2007/1205.pdf>

Goeree, Jacob. "Bidding for the future: signalling in auctions with an aftermarket." *Journal of Economic Theory* 108.2 (2003): 345-364.

Haile, Phillip A. "Auctions with private uncertainty and resale opportunities." *Journal of Economic Theory* 108.1 (2003): 72-110.

Hausch, Donald B. "Multi-Object Auctions: Sequential vs. Simultaneous Sales." *Management Science* 32.12: 1599-1610.

Holt, Charles et al. "Auction Design for Selling CO2 Emission Allowances Under the Regional Greenhouse Gas Initiative." Oct 2007. 5 Dec 2008 <http://www.rggi.org/docs/rggi_auction_final.pdf>.

Joskow, Paul L., Richard Schmalensee, and Elizabeth M. Bailey. "The Market for Sulfur Dioxide Emissions." *The Amercian Economic Review* 88.4 (1998): 669-685.

Kagel, John H. "Auctions: A Survey of Experimental Research." *The Handbook of Experimental Economics*. Ed. John H. Kagel & Alvin E. Roth. Princeton: Princeton University Press, 1995: 501-586.

Krishna, Vijay. *Auction Theory*. San Diego: Academic Press, 2002.

Mezzetti, Claudio, Alexander Saša Pekeč, and Ilia Tsetlin. "Sequential vs. single-round uniform-price auctions." *Games and Economic Behavior* 62.2 (2008): 591-609.

Ortega Reichert, Armando. "A Sequential Game with Information Flow" in unpublished PhD thesis Models for Competitive Bidding Under Uncertainty (Stanford, Department of Operations Research 1968). *The Economic Theory of Auctions*. Ed. Paul Klemperer. Northampton, MA: Edward Elgar Publishing, Inc. 130-152.

Regional Greenhouse Gas Initiative. "Auction Notice for CO2 Allowance Auction 6 on December 2, 2009." 14 Jul 2009. 29 Jul 2009 <http://www.rggi.org/docs/Auction_Notice_Oct_06_2009.pdf>.

Regional Greenhouse Gas Initiative. "Auction Notice for CO2 Allowance Auction 3 on June 17, 2009." 13 April 2009. 29 Jul 2009 <http://www.rggi.org/about/history/old_auction_notices>.

Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams. "Bubbles, Crashes, and Endogeneous Expectations in Experimental Spot Asset Markets." *Econometrica* 56.5

(1988): 1119-1151.
US Congress. House. 111th Congress, 1st Session, HR 2454, American Clean Energy and Security Act of 2009 [passed US House; 26 June 2009]. 111th Cong., 1st sess. 8 Jul 2009 <http://thomas.loc.gov/cgi-bin/query/z?c111:H.R.2454:>.
United States Environmental Protection Agency. A-96-19 - Comments on $SO_2$ market. 6 Aug 1996.

# 10 APPENDIX 1 – Market Summary Plots

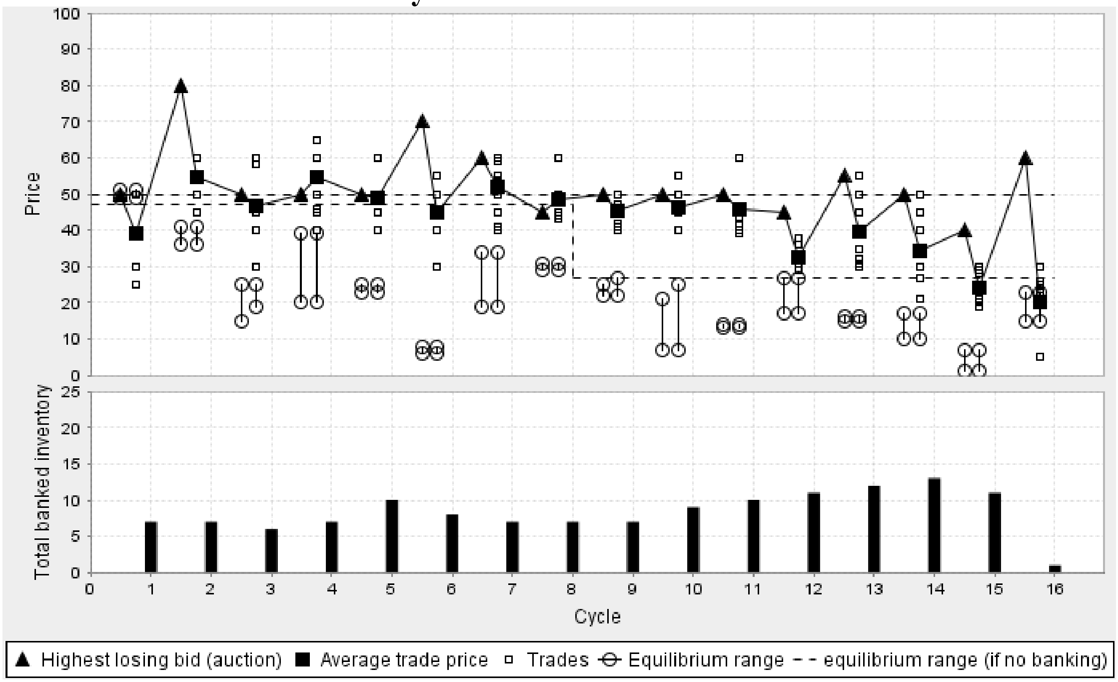


Figure 2: Uniform Public 25% auctioned - Run 1

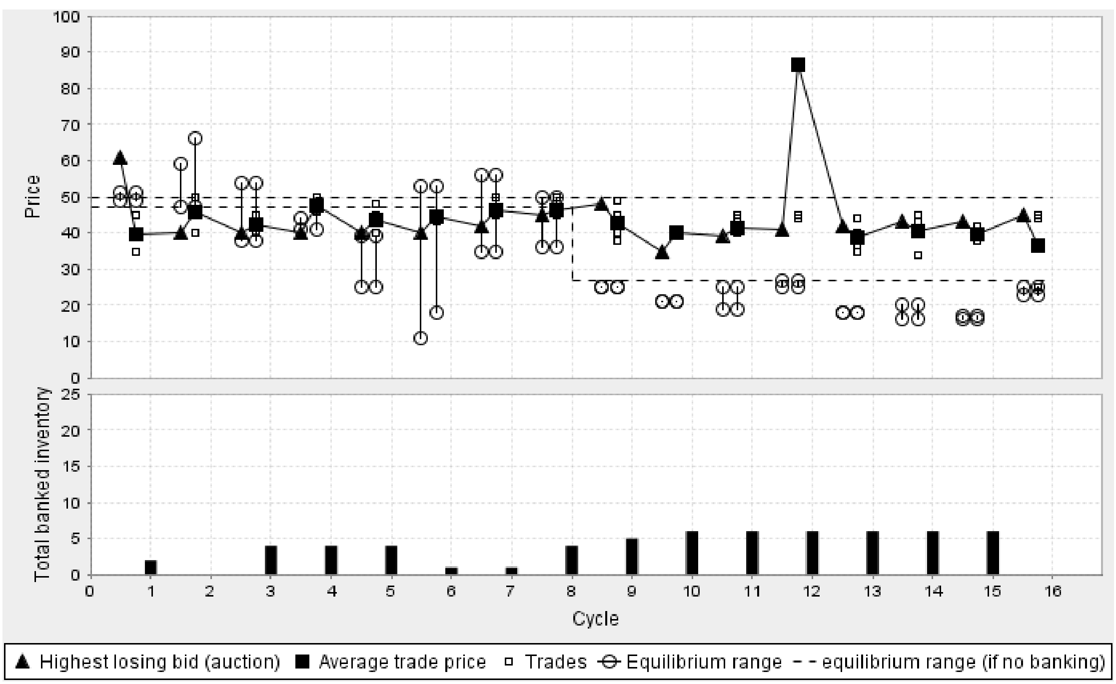


Figure 3: Uniform Public 25% auctioned - Run 2

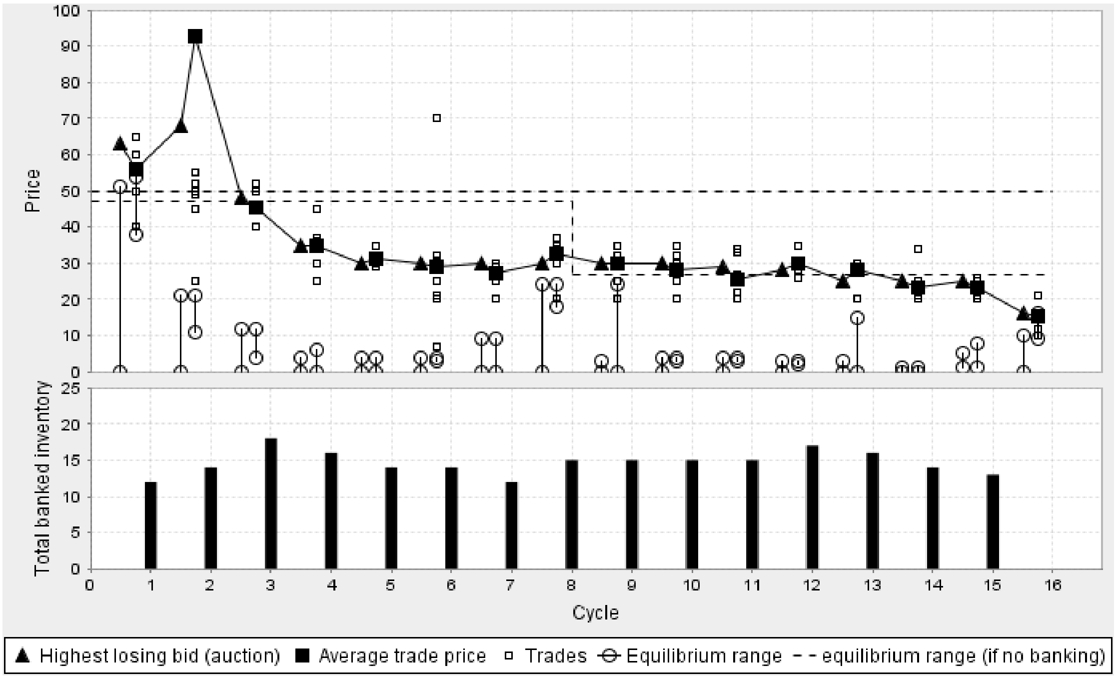


Figure 4: Uniform Public 100% auctioned - Run 1

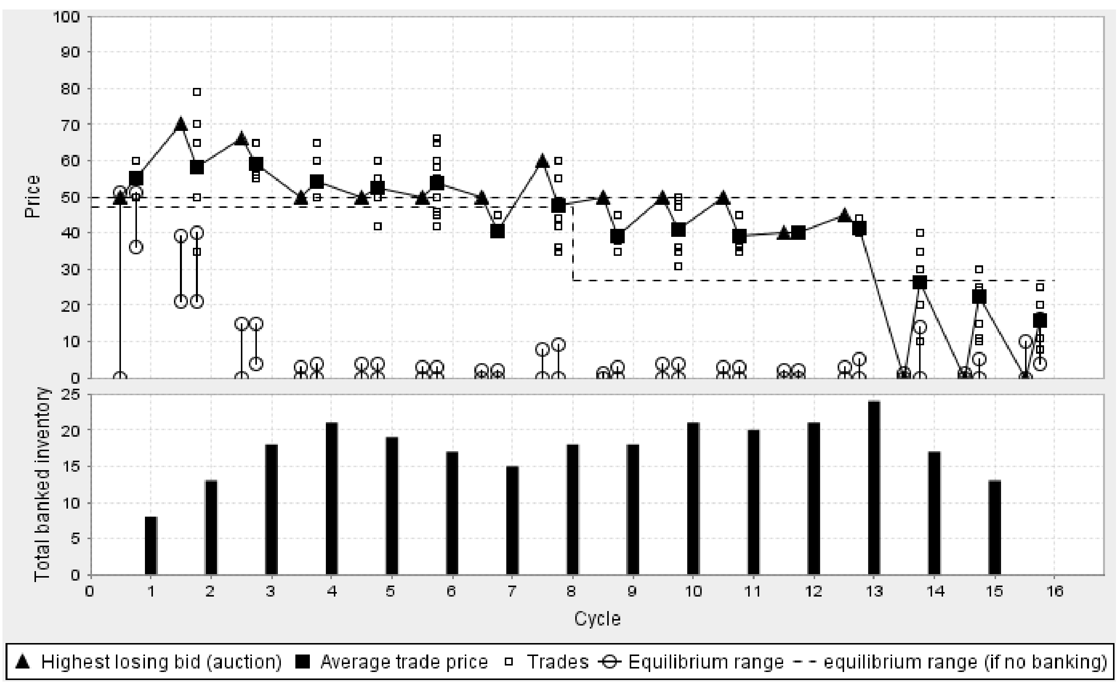


Figure 5: Uniform Public 100% auctioned - Run 2

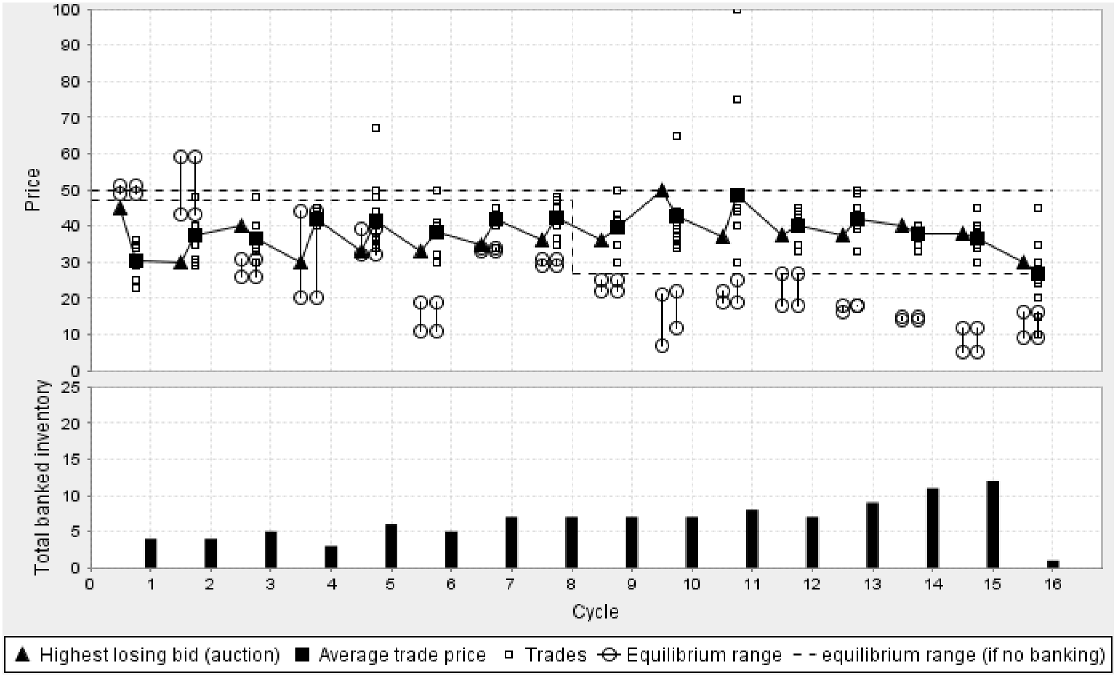


Figure 6: Uniform Private 25% auctioned - Run 1

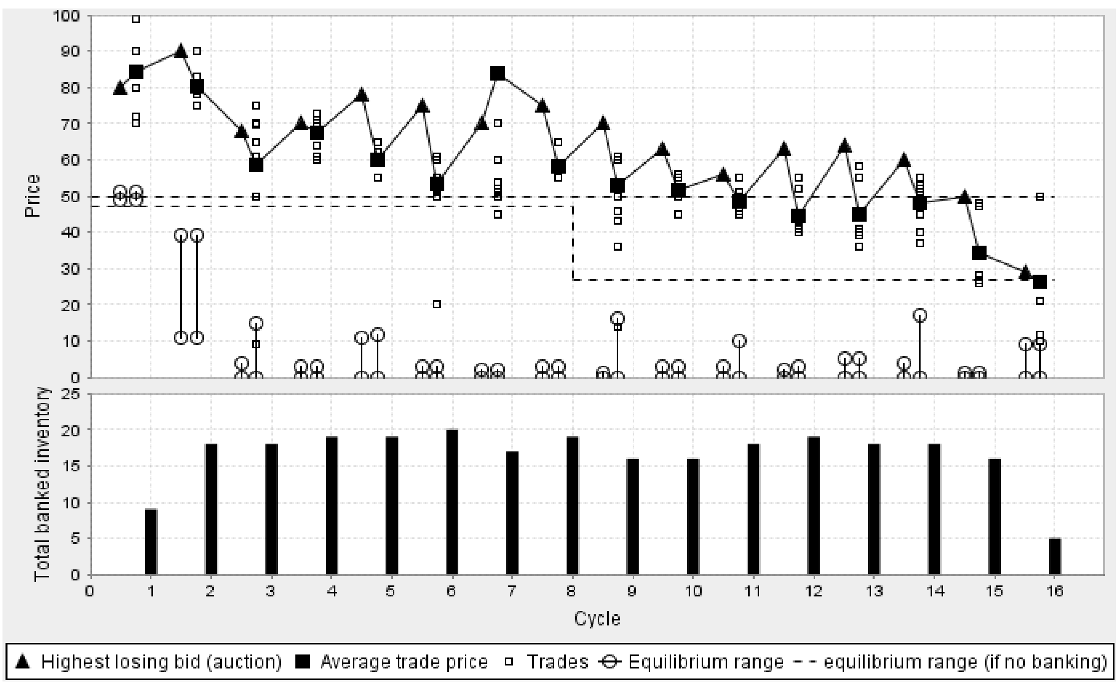


Figure 7: Uniform Private 25% auctioned - Run 2

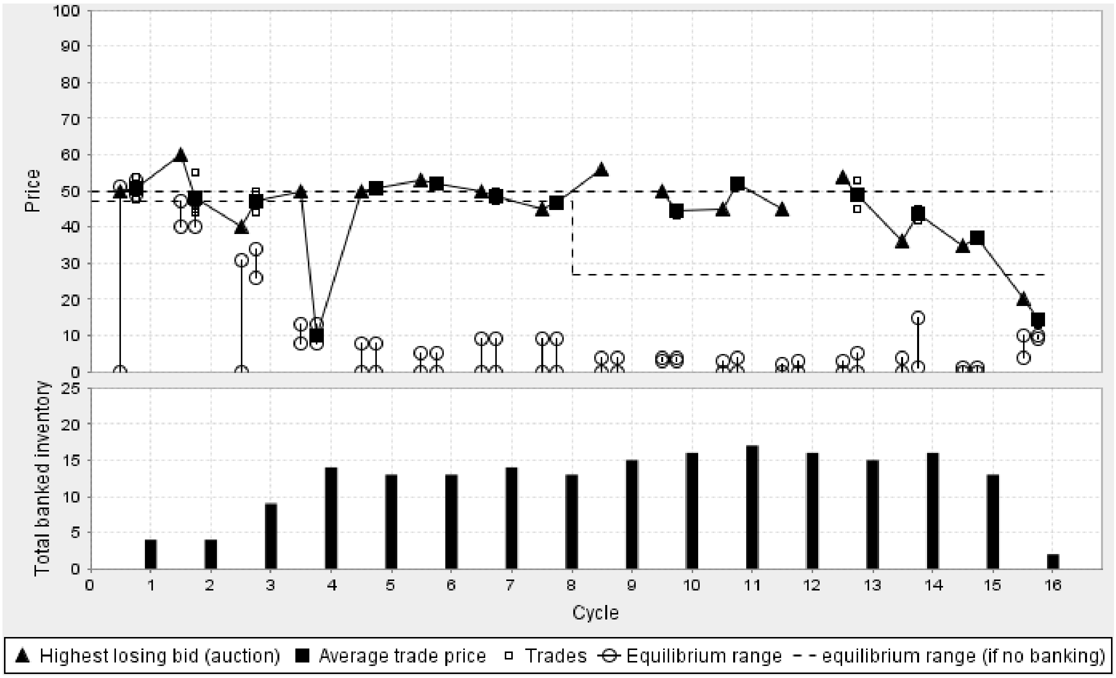


Figure 8: Uniform Private 100% auctioned - Run 1

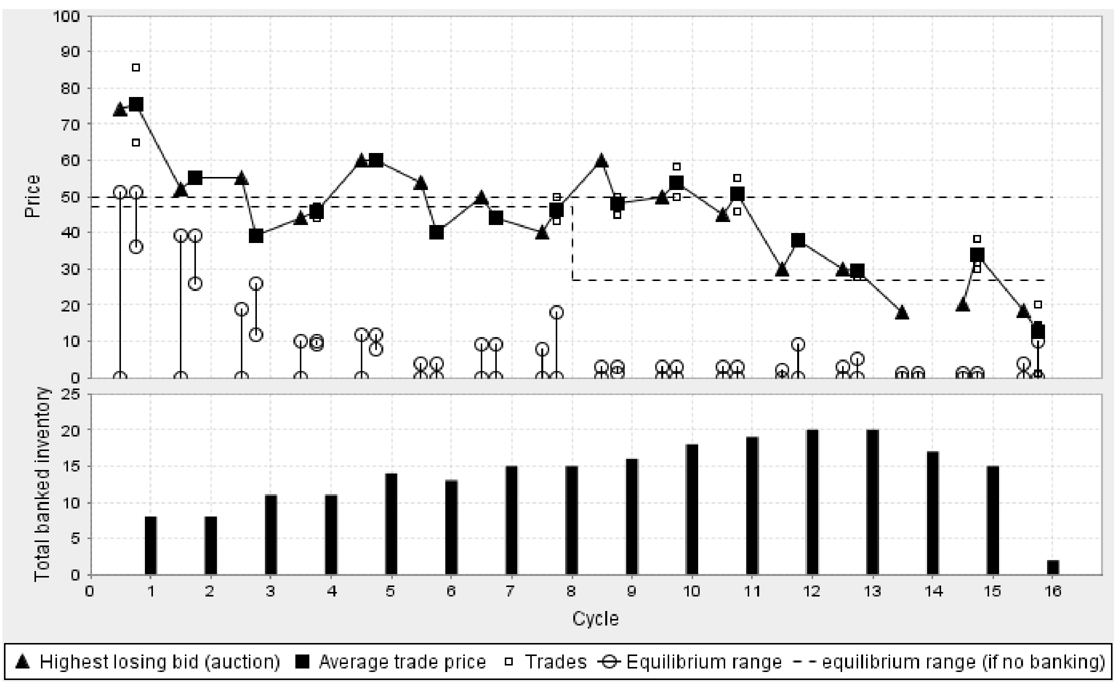


Figure 9: Uniform Private 100% auctioned - Run 2

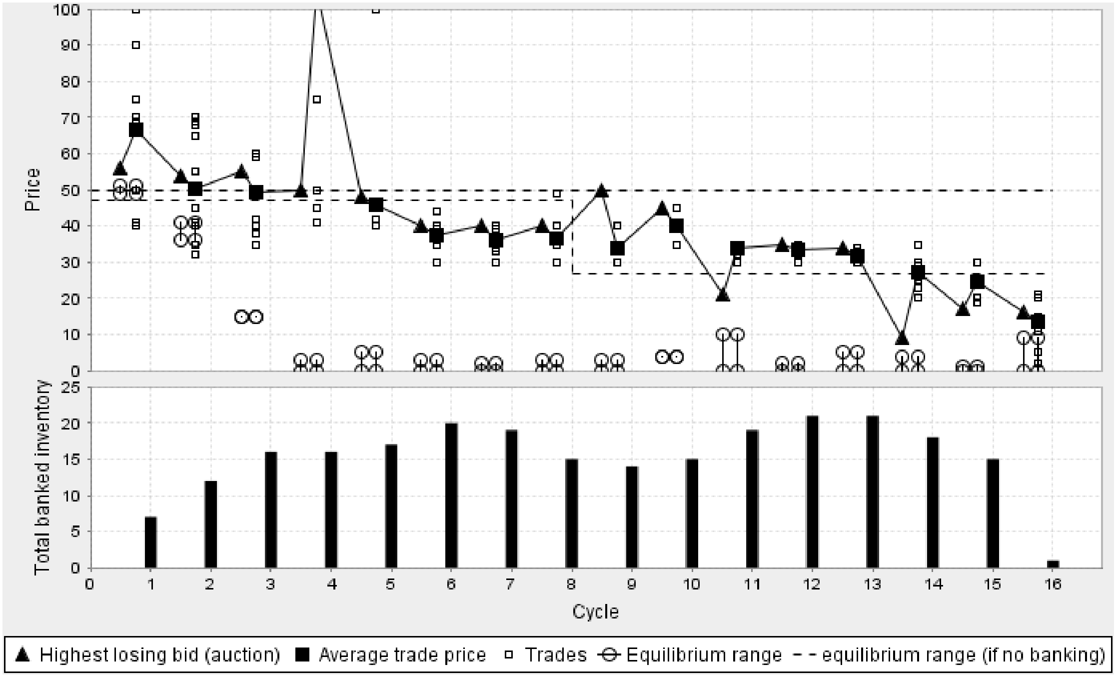


Figure 10: Discriminative Public 25% auctioned - Run 1

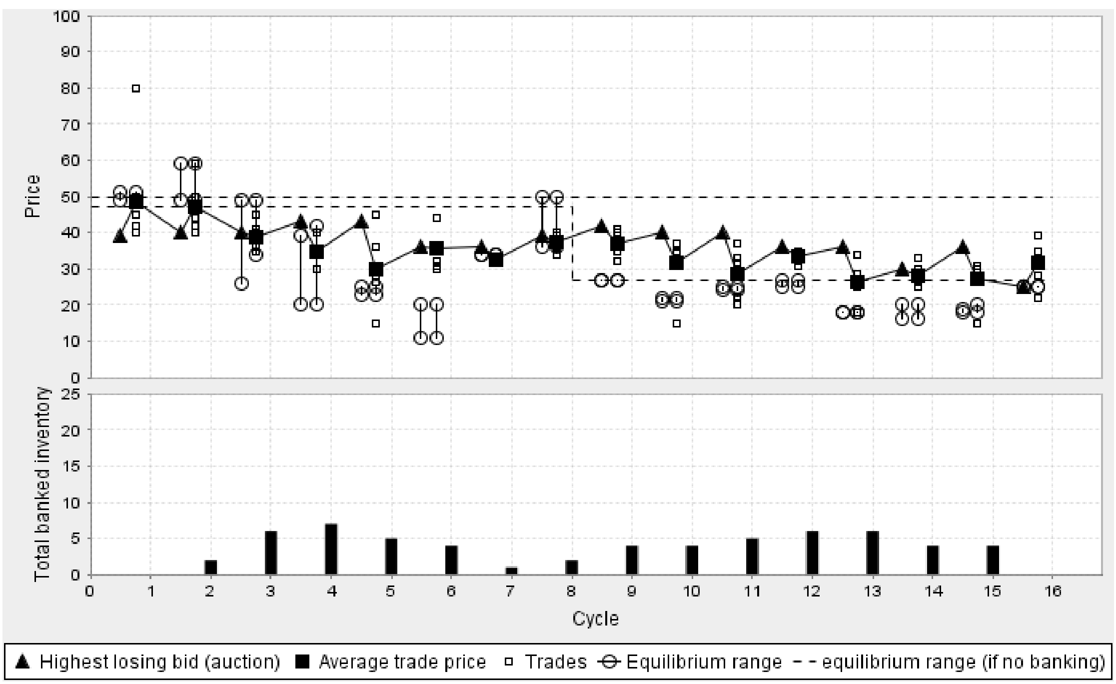


Figure 11: Discriminative Public 25% auctioned - Run 2

## 11 Appendix 2 – Instructions

Instructions for the experiment follow in this section.

Notation

As shown in section 3.5 of the paper, our experiment considered three two-level treatment variables:

Sealed bid design: uniform-price or discriminative-price sealed bid auction

Bid disclosure policy: complete bid disclosure (identity and price) or clearing bid only (price)

Fraction of allowance auctioned: 25% or 100%

Each modified the instructions given to the subjects. In what follows here, we include the text and figures for all treatments and values, and all changeable bits are highlighted in yellow. We note that subjects were given instructions specific to their treatment; they were not confronted with the choices shown here. Each highlighted portion begins with a bracketed identifier of the treatment to which it applies. For example, text presented to subjects in a market where 100% of the allowances were auctioned is prefixed by: [Variable: Fraction auctioned - Case: 100%].

**Instructions for repeated auction experiment – Live II Session**

## Introduction

This is an experiment in the economics of market decision making. The instructions are simple. If you follow them carefully and make good decisions, you may earn a considerable amount of money which will be paid to you in cash at the end of today's session. This is your final session for this experiment.

Please do not talk to your neighbors during this session, and do not share any private information.

In this market, all transactions are conducted in terms of *experimental dollars* (e$). You have been given a starting capital credit of e$1000. At the end of today's session, you will be paid in cash at a rate of 100 experimental dollars to 1 real dollar. [Variable: Fraction auctioned - Case: 100%] You will each also receive an additional $10 to complete today's session.

## Market cycles

In this market, you will participate in a sequence of market cycles. Today's market consists of 16 market cycles. A single market cycle consists of four steps:

1. Information Update
2. Auction
3. Trading
4. Production

Later in these instructions, I will explain what happens in each step.

In today's first cycle, I will stop the market before each step and read instructions specific to that step. Also during the first cycle, we will pause after each step and you will be given a sheet of exercises. These exercises are not graded; they are only meant to ensure that you each clearly understand the rules of the market.

In the remaining cycles, I will not read any additional instructions, but I will be happy to answer any questions that you may have. At any point today, if you have a question, ask.

## Input and output units

In this market you acquire and trade *input units*. You generate money in this market by using your input units to produce *output units*. It costs one input unit to produce one output unit. You can also buy and sell input units with the other market participants.

You profit in this market when you can acquire an input unit at a price less than either:

1. The amount you subsequently earn by using the input unit to produce an output unit

   - OR -

2. The price at which you subsequently sell the input unit to another player.

I will describe the mechanics of buying, selling, and producing later.

Input units

In each market cycle you can purchase input units in the auction, and you can purchase input units from the other players during the trading phase. Of course, this means you can also sell input units during the trading phase.

[Variable: Fraction auctioned - Case: 25%] Additionally, some of you will receive a fixed quantity of input units at the beginning of each cycle in the information update phase. The quantity you receive is a function of your username which has been randomly assigned to you.

Your inventory of input units is private information. No player can see any other player's inventory of input units. Do not share this information.

In the Production phase of each cycle, you can use up to five input units (assuming you have them in inventory) to produce output units. When you produce an output unit from an input unit, the input unit is removed from your inventory.

Input units that are not used to make output units are carried over into the next round. They do not expire. However, input units in your inventory after the final production phase (in cycle 16 today) are wasted. They have no value.

Output units

The money you generate from producing an output unit is shown in the table at the top left of your screen. The first column shows the order of production. The middle column shows the value of producing each individual output unit. The right column shows the total cumulative value of producing multiple output units.

Your table of values is five entries long, and the middle column (value of a single output unit) is descending in value. This means that the first output unit you produce is the most valuable. The second in the second most valuable, and so on. This is true for all players in the market.

Consider the following example.

| Production order | e$ Unit value | e$ Cumulative value |
|---|---|---|
| 1st | 86 | 86 |
| 2nd | 84 | 170 |
| 3rd | 41 | 211 |
| 4th | 27 | 238 |
| 5th | 18 | 256 |

The first output unit you produce in this cycle is worth e$86 to you. The second is worth e$84, so if you can produce two output units, you will increase your winnings by e$170 (86+84). If you produce four, you will increase your winnings by e$238 (86+84+41+27).

Your value table is different from everyone else's. No player can see any other player's table. This information is private. Do not share it with anyone.

Profiting with input units and output units

The auction and the trading phases provide you the opportunity to purchase input units. You can profit by buying an input unit for less than you can eventually earn from it (by producing an output unit or by selling it to someone else).

*Production profit*

If you produce an output unit from an input unit, your profit is given by

production profit = production value – original acquisition price

For example, assume you have e$100 and no inventory. If you purchase an input unit for e$50, you will be able to produce one output unit to generate e$86 of revenue and, therefore, e$36 (86 – 50 = 36) of profit.

*Sale profit*

If you sell an input unit (rather than produce output), your profit is given by

sales profit = sale price – original purchase price

For example, assume you originally purchased an input unit for e$20. If you sell it later for e$22, you will have earned a profit of e$2 (22-20).

*Balancing production and sales profit*

For each new input unit added to your inventory, you will choose whether to produce an output unit or to sell the input unit to another participant.

For example, assume you are still using the example value table above. Also, assume you have one input unit inventory. In the trading phase, you have the opportunity to sell the input unit for $50. The value table, however, tells you that your first output unit this cycle is worth e$86. Regardless of your original payment price for the input unit, the production profit is $e36 (e$86-e$50) larger in this case.

For another example, assume you have three units in inventory. In the trading phase, you are again given the opportunity to sell at $50. The value table tells you that your third output unit this cycle is worth $41. In this case, the sale profit is e$9 (e$50-e$41) larger than the production profit, regardless of your original payment price.

*Which units am I selling and buying?*

When you purchase a new input unit, it goes to the bottom of your value list. When you are selling input units, you are selling the input unit at the bottom of your value list.

For example, assume you have the example value table and four input units in inventory. If you purchase an additional input unit, its production value is e$18 to you. If you sell a unit and reduce your inventory to three input units, your production value of the input unit you sold would have been e$27.

When you are producing output, the first unit you produce in a cycle is the most valuable (the first row of the value table). The second is the second row, and so on.

**How the market works**

Each of the market cycles includes four phases: information update, an auction, trading, and production.

[Variable: Fraction auctioned - Case: 25%] In each market cycle, a fixed number of input units are released into the market. Some are awarded from existing agreements; others are auctioned.

– OR –

[Variable: Fraction auctioned - Case: 100%] In each cycle, a fixed number of input units are added to the market with an auction.

The first step, information update, is on your screen presently.

Information update

You do not need to do anything at this market step, but your information does change.

At the beginning of a new market cycle, you will all be given a new set of production values (the values in the table in the top left of your screen will change).

[Variable: Fraction auctioned - Case: 25%] Additionally, some of you have existing agreements that give you a number of input units at the beginning of each cycle. You do not have to pay for these input units. This quantity is fixed, and you will receive this amount every time a new cycle begins. You may see the number of input units that everyone else receives each cycle by clicking the „View Other Players' button at the bottom of your screen.

Auction

The auction is your first opportunity to buy input units each cycle. The number auctioned each cycle is constant and is shown at the top of your screen. All of you may participate in the auction.

The auction allows you to increase your inventory of input units. The auction does not allow you to sell anything from your inventory, nor does it allow you to produce output units.

Your screen will allow you to submit a series of bids in the form of a demand curve. Do not submit your curve until you have decided on all of your bids. Your demand curve can be of any length from zero to the number of input units offered in the auction. Each bid in your demand curve must have a price greater than e$0. The market ignores all bids of e$0.

Once you have all submitted your demand curves, each curve is separated into a set of individual bids, each for one input unit, and the complete set of bids is sorted from highest to lowest. If two bids are for the same price, then the offer that was submitted earlier is set higher on the list. Input units are then sold to the owners of the highest bids. [Variable: Auction design - Case: uniform price] The lowest price bid that results in a sale [Variable: Auction design - Case: Discriminative price] The highest price bid that does not result in a sale is called the *clearing bid*. All bids with prices above this bid yield sales; those below do not. [Variable: Auction design - Case: uniform price] The sale price for each input unit will be the price of the clearing bid. [Variable: Auction design - Case: Discriminative price]The sale price for each input unit will be the price of the bid itself.

[Variable: Bid disclosure - Case: Full disclosure] The results of this auction will be made public to all participants afterward. You will be able to see the clearing price of the market, and you will see all information about all of the submitted bids (price and bidder identity). Note that this means everyone can see your bid information, as well.

[Variable: Bid disclosure - Case: Clearing bid only] After the auction, you will be given the clearing price and quantity. You will be able to see your bids, but you will have no more information about the other bids in the auction.

To best illustrate the rules, consider an auction in which four input units are being sold. Three bidders participate: Larry, Curly, and Moe. Their demand curves are shown. You will be submitting curves of this form, but you will choose the prices and the number of items being auctioned is different.

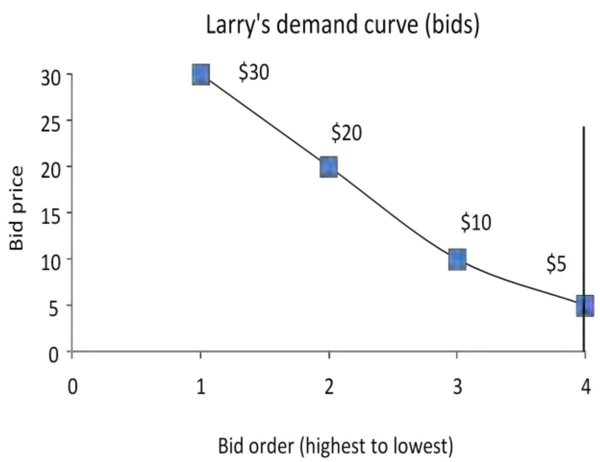


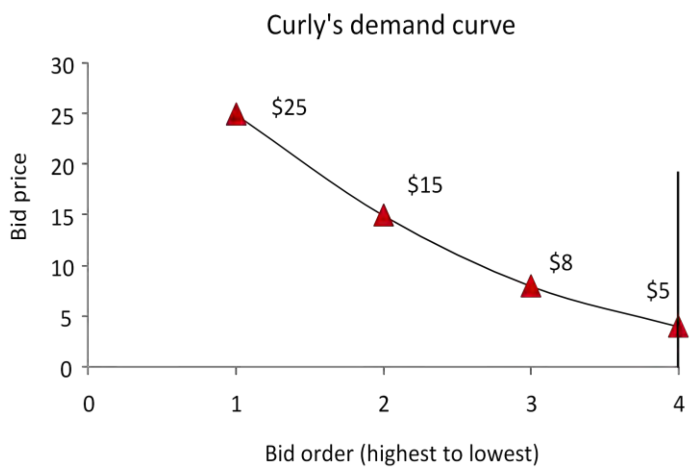



You can see that both Larry and Curly have submitted bids for the full number of inputs being auctioned. Moe has not. Your demand curves can be of any length up to the number of input units being auctioned. You do not have to submit a curve if you choose not to do so.

Once the demand curves have all been received, they are merged, and the individual bids are sorted highest to lowest. The following plot shows the whole market demand curve that results from the three demand curves.

[Variable: Auction design - Case: uniform price]

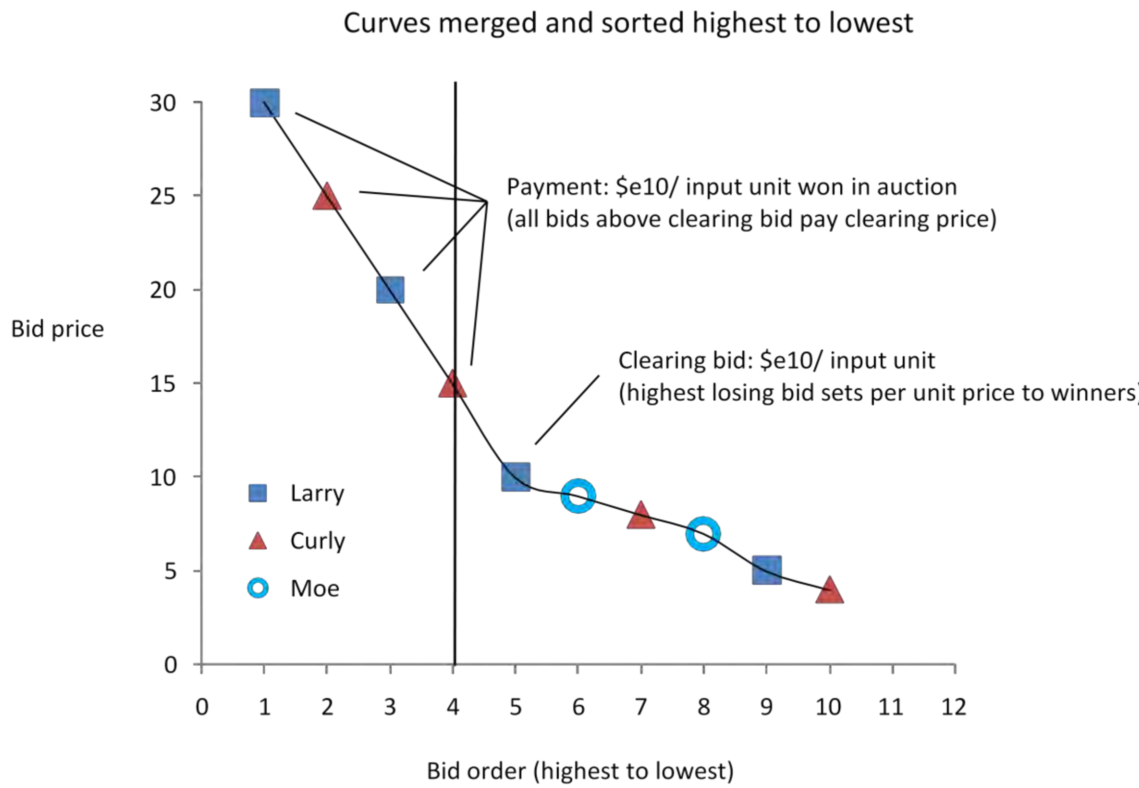

[Variable: Auction design - Case: Discriminative price]

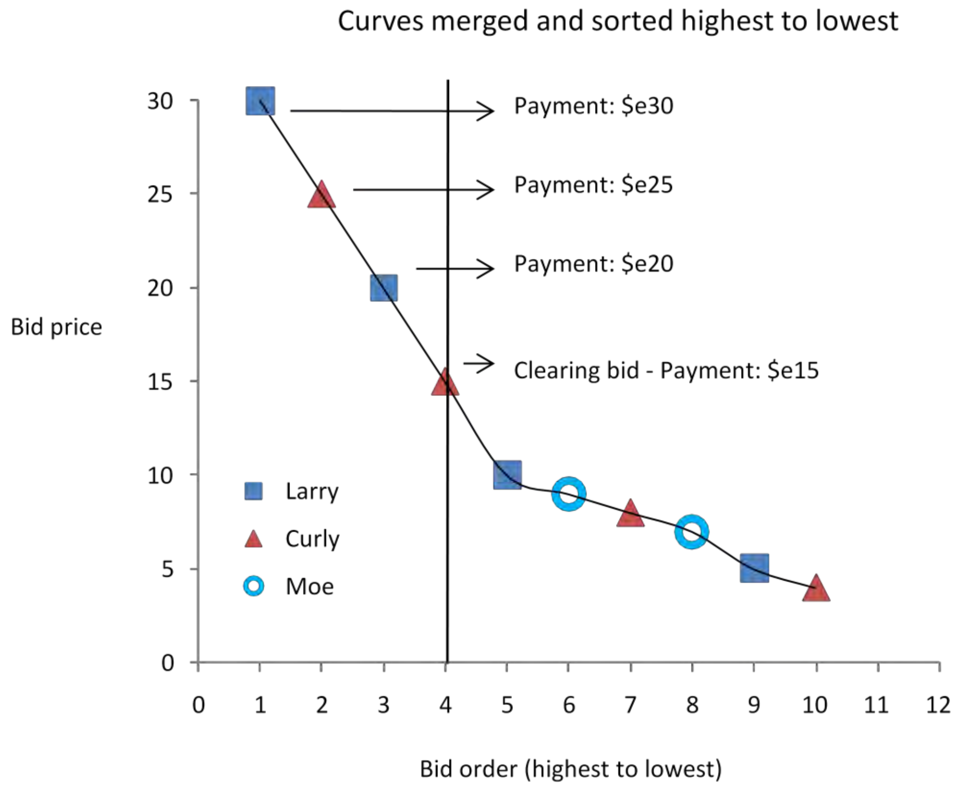


You can see that Larry and Curly have the four highest bids in the auction (as well as some losing bids each). Moe's bids are all below the clearing bid. Larry and Curly are each awarded two input units.

[Variable: Auction design - Case: Uniform price] Remember, the highest losing bid sets the per unit purchase price for all winning bids. Notice that since the clearing bid sets the price for all sales, if you bid higher than the clearing bid, you will pay less than your actual bid. Here, Larry's third bid for e$10 is the highest losing bid. Therefore, the four winning bids are awarded input units at e$10 each. Larry wins two input units and pays a total of e$20. Curly also wins two input units and pays a total of e$20. Moe wins nothing and so pays nothing.

[Variable: Auction design - Case: Discriminative price] Remember, each winning bid results in a payment for the value of the individual bid. Since you will pay your bid price if you win an input unit in the auction, you have complete control over the price you pay. However, if your bid is too low, you may not win any input units. Here, Larry pays a total of $50 ($30 + $20) for his two input units. Curly pays a total of $40 ($25 + $15) for his two input units. Bids below the clearing bid are not awarded input units and do not result in payment.

Note that in this experiment, [Variable: Bid disclosure - Case: Full disclosure] because all bid information is public after the auction, you will be able to see this plot in the market history control on the right of your screen. [Variable: Bid disclosure - Case: Clearing bid only] because all bid information is private, after the auction, you will not see the whole market curve above. Rather, you will see the clearing bid and your own bids -- nothing more.

Now, you can see in this example that the submitted demand curves are themselves sorted highest to lowest. Because of the way winning bidders are identified in the auction, your curves must also be sorted highest to lowest. The software will reject any demand curve that does not satisfy this rule.

Finally, you may not go into debt purchasing input units. The software will prevent you from submitting a collection of bids that are greater in cost than your current experimental winnings.

Today, you will have at most two minutes to submit your bids in each auction.

<Break to conduct auction>

Now that the auction is concluded, you can see the results in the „Market history' tree on the right side of your screen. Please take a moment to view the results of the auction.

<u>Trading</u>

The trading phase provides an opportunity to buy and sell input units. You can profit by selling input units to other players for more than you can earn by using them to produce output units, or you can increase your inventory of input units by purchasing them from the other players.

There is no limit to the number of offers or the number of sales executed during this phase.

The software allows you to post buy offers and sell offers to the market. When the highest buy offer is greater than or equal to the lowest sell offer, a sale is made. The price of the sale is determined by the offer (buy or sell) that is submitted second.

To illustrate, Curly offers to sell an input unit for e$20. Moe offers to buy for e$10. No sale occurs. Larry then offers to buy for e$20. The result is that Larry buys a unit from Curly at e$20.

To illustrate again, Curly again offers to sell for e$20. Moe then offers to buy for e$30. The result is that Moe buys a unit from Curly at e$30. The sale occurred because the buy offer was greater than the sell offer. The sale was executed at $30 because the second offer was for $30. Had Curly offered to sell at $20 after Moe offered by buy at $30, the sale price would have been $20.

When submitting buy offers, you must submit an offer that is higher than the standing market buy offer. When submitting sell offers, you must submit an offer than is lower than the standing market sell offer. When the market first starts or when a sale has just been made, there are no restrictions on the first sell or buy offer.

As with the auction, you may not go into debt to purchase input units. Also, you may not sell input units if you do not have any in inventory. In both cases, the software will prevent you from doing this.

In each cycle, trading will last two minutes.

<Break to trade>

Now that the trading phase is complete, you can see the record of trades in the history panel on the right of your screen.

<u>Production</u>
The production phase allows you to spend input units to produce output units and, therefore, increase your winnings. When you spend an input unit, it is removed from your inventory.

The table at the top of your screen gives the value of producing output units. Remember, one input unit is required to generate one output unit.

For example, in the following table, if you produce two output units, you will earn e$170. If you produce four output units, you will earn e$238.

| Production order | e$ Unit value | e$ Cumulative value |
|---|---|---|
| 1st | 86 | 86 |
| 2nd | 84 | 170 |
| 3rd | 41 | 211 |
| 4th | 27 | 238 |
| 5th | 18 | 256 |

If you choose not to use all of your input units in a particular cycle, they will carry over into the next cycle. They do not expire. However, input units remaining in inventory after the last cycle (cycle 16) have no redemption value.

In each cycle, you will have a maximum of one minute to select your production level.

<Break to produce>

SUMMARY

You have now completed a full market cycle. Fifteen more remain in today's event.

To close today's instructions, a couple of points are highlighted.

1. All information and transactions in the experiment are conducted in terms of experimental dollars.
2. You can profit by producing output.

   production profit = production value – original acquisition price

3. You can profit by selling input units to other players.

   sale profit = sale price – original acquisition cost
4. For a single input unit, you can either produce an output with it or sell it to another player. You cannot do both.
5. Input items remaining in your inventory after the last production phase are wasted. They have no value.
6. In each auction, [Variable: Auction design - Case: Discriminative price] you pay your full bid price for each of your bids that results in a sale. [Variable: Auction design - Case: Uniform price] you pay the clearing price for each of your bids that results in a sale.
7. After each auction, your bid information is [Variable: Bid disclosure - Case: Full disclosure] PUBLIC. All of you will be able to see every bid submitted in the auction. [Variable: Bid disclosure - Case: Full disclosure] PRIVATE. No other player will ever see your bids.
8. Please do not talk to any other participant or reveal any of your information during the experiment. If you ever have a question, please raise your hand, and I will be happy to answer it.

   Are there any questions now? Feel free to ask questions throughout today's session.