Auction Design for Selling CO_2 Emissions Allowances Under the Regional Greenhouse Gas Initiative. Holt et al. October, 2007

Addendum: Response to Selected Comments

Charles Holt, William Shobe, Dallas Burtraw, Karen Palmer, Jacob Goeree, Erica Myers

April 8 2008

Summary: A final report summarizing research on auction design for RGGI was submitted to the New York State Energy Research and Development Authority (NYSERDA) in November 2007. Subsequently several constructive comments and suggestions have been made by interested parties. This addendum to the November 2007 report contains reactions to some of these comments, including some new findings based on research done subsequent to submission of the Final Report.

The main issues addressed in this addendum are (1) the choice of auction design (clock versus sealed bid) and how this affects the efficiency of the auction and the ability of parties to collude, (2) variations on design for clock auctions, and (3) auctions that combine different vintages of allowances. To summarize the main results:

- In laboratory auctions with communication among participants, successful collusion is more effective in clock auctions than in discriminatory and uniform price auctions.
- An analysis of the 'chat' (instant message communications between bidders prior to submitting bids) indicates that clock auctions facilitate collusion by allowing bidders to focus on a single dimension (quantity reductions).
- The effects of this collusion are reflected in clock prices at or near reserve price levels, with subsequent trading at much higher prices in the spot markets.
- Results of our new experiments, conducted subsequent to the Final Report, indicate that the provision of information about the quantity of demand after *each round* of a clock auction does not improve price discovery of these auctions.
- This type of demand reduction in clock auctions echoes the striking results reported by Goeree, Offerman, and Sloof (2006) in a much simpler environment with full information provided about quantities demanded during the auction.
- The New England and New York ISO proposal that allowance owners be able to offer allowances for sale in the RGGI auction has definite advantages. We have different suggestions about how this might be implemented. The uncertainty of supply that results can help reduce the potential of collusion.
- Since RGGI allowances are bankable, a bid for a later vintage could be treated as a request to purchase *either* a later vintage *or* an earlier vintage, whichever is less expensive. Interpreting bids in this manner prevents a price inversion in which the uniform price for the later vintage is higher than the price for the earlier vintage, although theory suggests this price inversion is inefficient and would not occur in the secondary market. This addendum describes a simple procedure for combined vintage auctions that implements this idea.

(1) Choice of Auction Design

A central recommendation of our final report (November 2007) is to use a uniform price sealed bid auction in RGGI. One concern that emerged in our research was the stronger opportunity for collusion in clock auctions compared to sealed bid auctions. Below we discuss additional research on this question.

Effects of Explicit Collusion in Laboratory Clock Auctions

Our final report expresses some concern about the effects of collusion in clock auctions; these concerns were based on a reading of some related literature, but they were not directly supported by our own experiments. Subsequent to completion of the Final Report we conducted additional experiments to explore the issue further. We addressed the effects of collusion in clock auctions by running additional sessions with clock auctions in which participants could discuss any aspect of the auction in a chat room that was open prior to each round of bidding. The results provide strong evidence that collusion is more effective in a clock auction than under other auction formats. The average prices were lower for the clock format than for the uniform and discriminatory price auctions. Figure 1 shows average purchase prices for a series experimental sessions that each had a sequence of twelve auctions. The clock prices begin above $2.50 and decline progressively to levels that are well below the averages for the discriminatory and uniform-price treatments, which are essentially indistinguishable. For reference, the average Walrasian prediction (where supply equals demand) is shown as a thick dashed horizontal line just above $3.50.

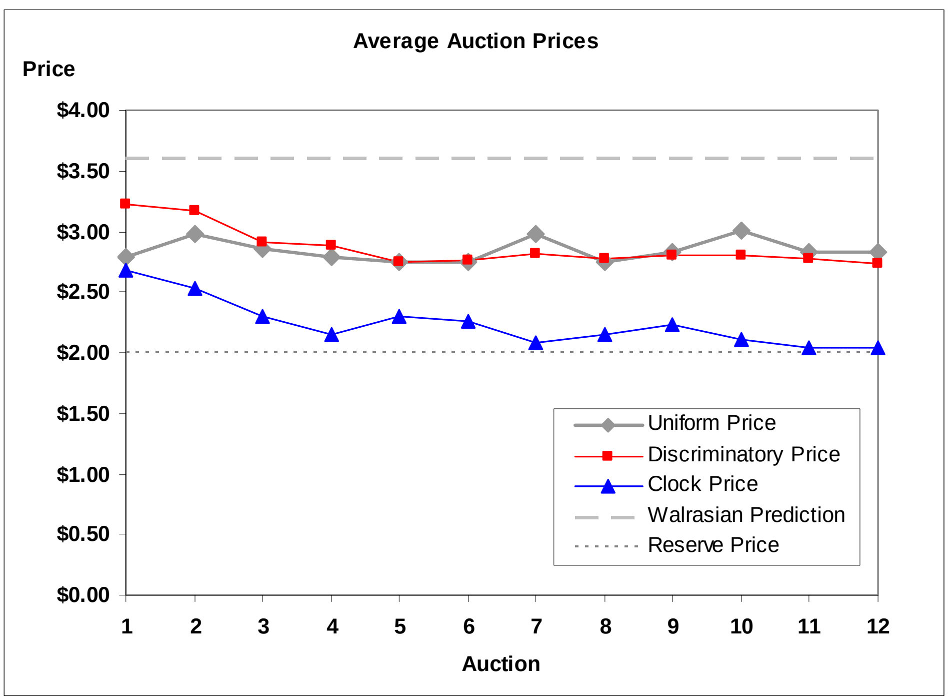

Figure 1: Evidence of Collusion Under Alternative Auction Formats

The averages shown in Figure 1 to some extent mask the sharply increased tendency for clock auctions with collusion to stop at the reserve price or at one or two bid increments above reserve, as shown in Figure 2. Two-thirds of the 36 clock auctions stop at the reserve price of $2.00, whereas the modal price outcome for the uniform price auctions is $2.90.

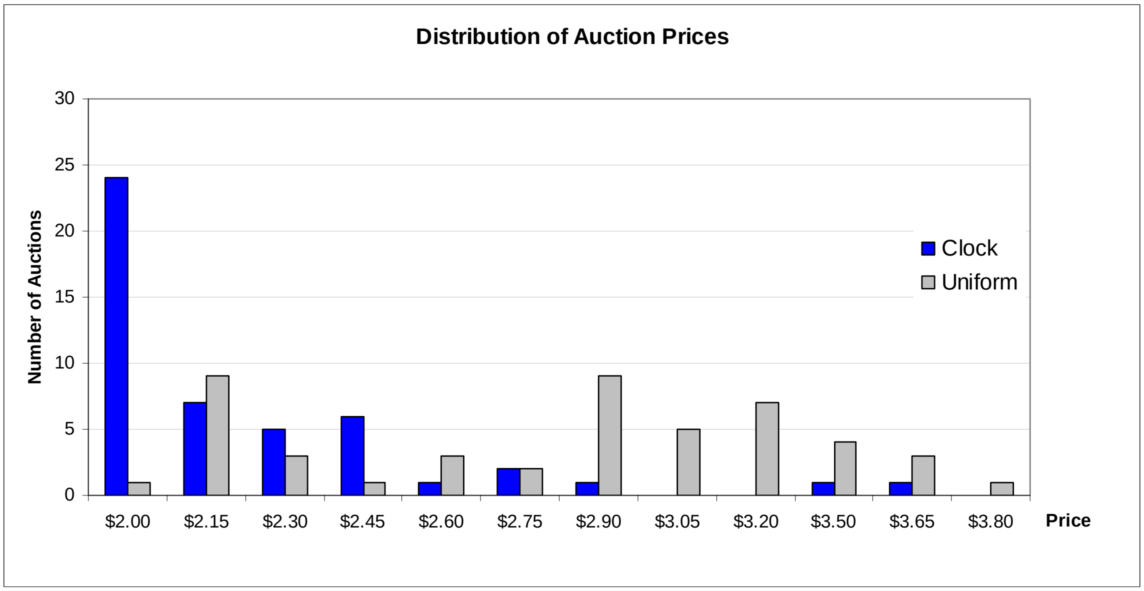

Figure 2: Price Distributions with Communication: Uniform vs. Clock

An analysis of the actual chat between subjects provides some insight for why collusion is more successful in clock auctions. Most of the initial proposals were based on suggesting quantity reductions for low and high users. The focus on quantity reduction in the clock auction sessions is revealed in some of the participants' comments: "again, bid for fewer permits earlier on so we can get permits cheaper" and "this will go 5X faster and will all make LOTS more money if everyone just cooperates the first time." One person suggested "so why doesn't everyone bid exactly the same amount that we ended last round [auction] with, since we keep getting the same clearing price." This plan permitted participants to obtain the same final permit allocation without the run up in prices that occurred previously. Of course, quantity discussions occurred with the other auction formats too, but the effect of the clock is to take out the price dimension so that bidders only have to reach an agreement in a single dimension, quantity.

One interesting feature of the data for the clock sessions with chat is that, while auction prices were typically near the reserve price, the subsequent trading in the spot market tended to be at much higher price levels that were closer to the Walrasian price, as indicated in figure 3. The bidders were defeating the clock auction with successful collusion, which typically results in inefficient allocations that were, to some extent, corrected by trading in the spot markets. This is why the aggregate efficiency results for the clock with collusion were not significantly different from the efficiencies for other formats, as shown on the right side of Figure 4.

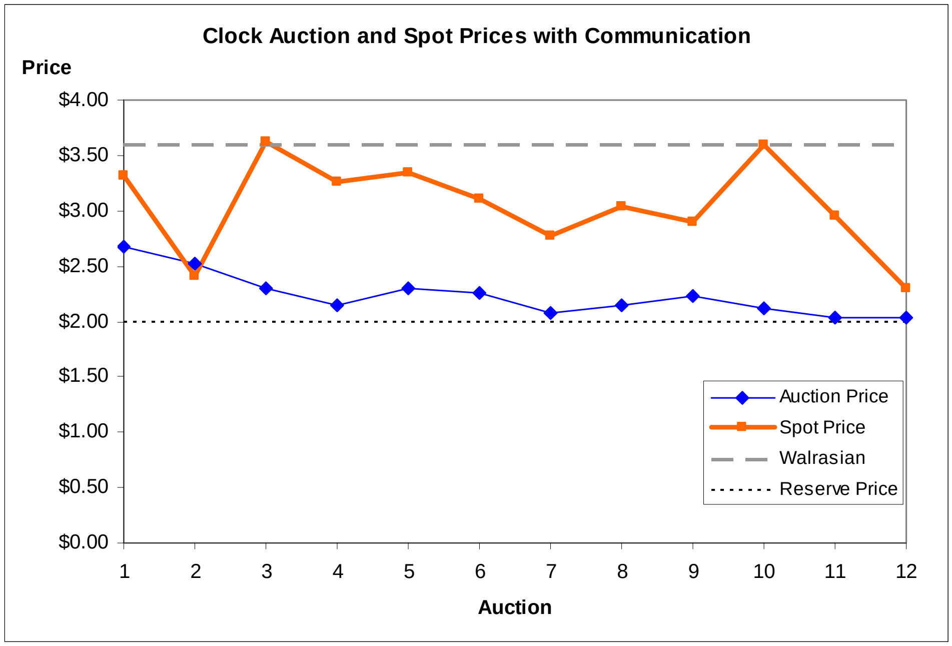

Figure 3. Collusion: Low Clock Auction Prices and High Spot Prices

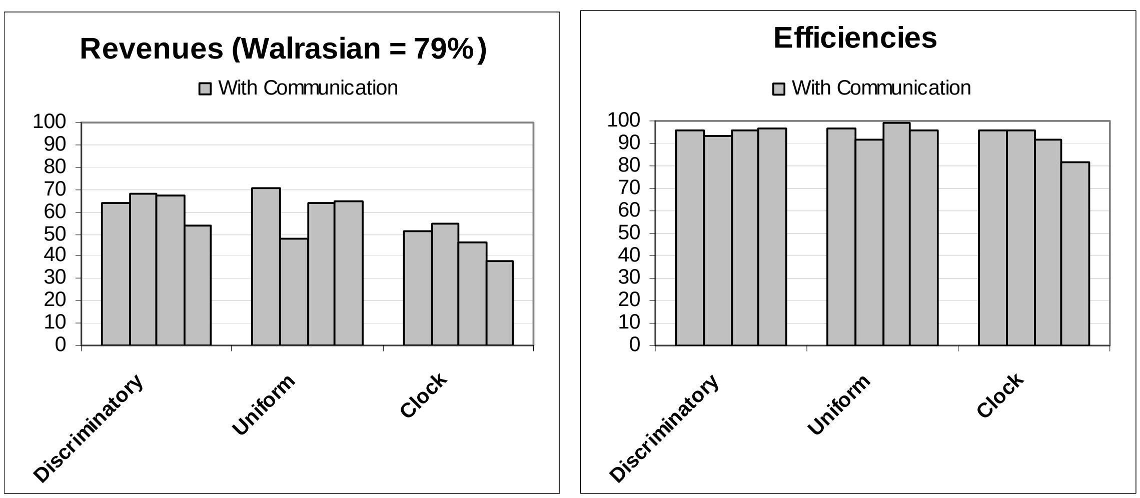

Figure 4: Collusion Sessions: Efficiencies and Revenues by Auction Type

There are several widely cited cases in which coordinated demand reductions are used to defeat ascending price auctions for broadcast spectrum. Just prior to the 2001 Austrian UMTS auction, Telekom Austria announced that it "… would be satisfied with just two out of the 12 blocks for offer and if the [five] other bidders behaved similarly, it should be possible to get the frequencies on sensible terms … but that it would bid on a third block if one of its rivals did …" The other bidders clearly understood and the bidding stopped after a couple of rounds, with each bidder obtaining 2 blocks (Klemperer, 2004, p. 136). The extreme symmetry of this situation differs from the bidding environment faced by bidders in RGGI auctions, but the success of collusion triggered by a public announcement is disturbing.

2) Variations on Design for Clock Auctions

Several comments suggested that the ideal design for a clock auction might have several features that we did not include in our early experiments. In subsequent work we have explored several additional design features.

End of Round Quantity Information

On average, the clock auction prices observed with collusion were quite close to the reserve levels, and therefore we did not think it worthwhile to redo these (very expensive) clock auction sessions with the additional feature of announcing the total demand after each separate round of bidding. With ex post quantity information after each round, bidders would be able to "signal" an intent to cooperate by reducing their demands to stop the clock early in one auction, in the hopes of inducing reciprocal cooperation from others in subsequent auctions. In fact, our empirical "sample of one" from the Virginia NOx auction conducted in 2004 suggests that end-of-round quantity information may even induce bidders acting individually to reduce quantities bid in order to stop a clock auction. In the NOx auction, bidders were not provided with end-of-round quantity information, since the George Mason design team (that included a subsequent Nobel-prize winning economist, Vernon Smith) advised that providing such information may encourage collusion. Late in the Virginia NOx auction, there was one bidder with a large share of the allowance requests; this bidder could have ended the auction by unilaterally reducing demand. As the clock price increased, this bidder began shading requests by relatively small amounts, in an apparent attempt to "feel around" for the edge needed to stop the auction. We believe that this bidder would have ended the auction sooner, with a quite modest "demand reduction," if the needed information about demand "overhang" had been available.

The decision of how much information to provide during the auction should be affected by opinions and evidence about the tradeoff between the effects of additional collusion facilitated by the information and the added ability of the auction to track the market clearing price as a result of the added information. In order to assess this tradeoff, we ran additional sessions with provision of this information on a round-by-round basis. This "price discovery" treatment involves a series of 6 auctions, with an unannounced demand shift (caused by a change in fuel/input costs) that produces a large increase in the Walrasian price for the final 3 auctions. Figure 5 shows the clock price averages for the series of auctions, again with the Walrasian price predictions shown as a dashed line. On average, the ability of the clock auction to track the demand shift is not improved by the provision of end-of-round quantity information. The reason is apparent from the strong tacit collusion that generated demand reduction and sharply reduced clock prices in some cases.

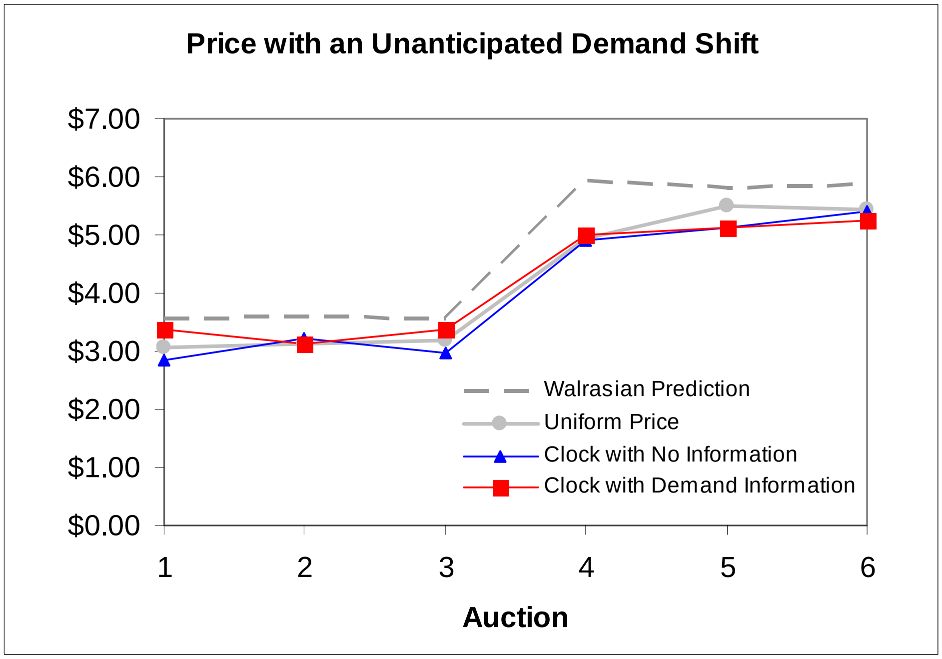

Figure 5: Tracking an Unanticipated Demand Shift:
Clock Auctions with and without Excess Demand Information

Goeree, Offerman, and Sloof (2006) report a multi-unit auction experiment designed to assess the likelihood of demand withholding in a very simple environment with no opportunities for communication. In their experiment, groups of 3 participants were competing for 6 available units. Each bidder received a randomly determined private value for three units, so the total demand at a zero price could be as high as 9. The experiments were either run as discriminatory auctions or as clock auctions in which bidders were informed when anyone reduced their bid quantities. Tacitly coordinated "demand reduction" tended to stop the clock at low prices, and auction revenues were much higher in the discriminatory auction, i.e. 151 versus 40 for the clock, a difference that was statistically significant at the 1% level. In addition, there was less revenue variability across auctions in the discriminatory treatment. The revenue advantage of the discriminatory auction even persisted in other asymmetric treatments in which two of the bidders in each group were "incumbents" with strong incentives to bid high and exclude the third bidder ("entrant"). This paper did not consider a uniform price sealed bid auction, but the authors summarize the literature relevant to this comparison: "Pooling the results from these different studies suggests that demand reduction is more pronounced in ascending auctions than in uniform-price auctions." (Goeree, Offerman, and Sloof, 2006, p. 2.)

Intra-round Bidding Schedules

One concern about clock auctions that is sometimes raised is that clock auctions require greater time and resources on the part of bidders. At least one large commercial auction service vendor has noted that its clients object to the amount of time it takes to run a clock auction and to the fact that the termination time is not fixed. They noted that energy prices change hourly in some cases. We are not in a position to evaluate these claims, which they provided as an argument in favor of the "hard close" bidding formats that they typically use for clients in energy procurement auctions.

One way to make a clock auction run more quickly is to use larger bid increments for each "tick" of the clock. Doing so may reduce the efficiency of the auction, but the issue can be addressed by letting bidders submit a second set of bids within each round of bidding in a clock auction. These intra-round bid schedules are quantities demanded at prices between the clock price for the previous round and the clock price for the upcoming round. The resulting "fine bid increments" between rounds would be used if not all allowances sell at the next higher clock price. Course bid increments may reduce the efficiency of a clock auction, but the same is true for other auction formats. In the experiments, we used the same bid increments for all formats in a given treatment.

We view the use of coarser price increments and intra-round bidding schedules with fine increments as a plausible "fix" for another problem specific to the clock, delay and uncertainty about ending times, This fix does require two sets of bids per round, and submitting a whole "demand function" in each round adds complexity. Would this procedure work better than the alternative of having a regular clock with bid increments half as large and rounds half as long, but with only one set of bids per round? The answer is not known. We continue to believe that laboratory tests of alternative formats are most informative when the bid increments are held constant across formats.

Exit Bids

Another fix for the clock's delay and uncertain ending problems is to allow bidders to submit a single demand function bid (quantities requested at a whole range of possible prices) at the start of the auction, instead of updating their bid in each round. This form of bidding is known as submitting *exit bids*. Then this demand function would be used to determine bidders' requests at each stage of the clock auction. If all bidders selected this form of bidding, the result would be equivalent to a uniform price, sealed-bid auction. So this exit-bid option has some appeal in terms of convenience. Note, however, that using this option prevents bidders from learning and responding to information revealed during the auction. Any "price discovery" benefits of using a multi-round auction would be diminished to the extent to which bidders adopt exit bids. Therefore, we do not believe that allowing exit bids would improve the performance of clock auctions in comparison with other formats used in the laboratory tests. An intermediate possibility might be to allow exit bids as a proxy strategy, but to allow bidders to withdraw and revise those bids if they choose to monitor the auction. Bidders may chose to monitor intermittently, and may chose to act only if the auction is performing in a way that is significantly different from expectations.

Supply Function Bidding

One interesting suggestion made by Peter Cramton in the comments he prepared for the New England and New York ISO is that bidders be given the option to sell (as well as buy) allowances in the auction (Cramton, 2007). This could be accomplished by letting them submit supply functions, i.e. quantities and associated minimum sale prices. His recommendation is that the resulting aggregate supply function be made public prior to the auction. Sales possibilities, for example, are incorporated into the SO_2 auctions run by U.S. EPA. We like the idea of allowing sales as well as purchases in the RGGI auctions, in part because such sales could provide more of a quantity cushion. Our initial reaction is that it might be best to allow sales offers to be submitted simultaneously with all other bids in a uniform price auction. In this case, the bids generate a demand curve, the sales offers generate a supply curve, and the clearing price could be determined by the intersection. This procedure would reduce to the recommended uniform price auction in the event that no sales offers were received. With sales offers, this type of market is known as a "call market," which is commonly used in electronic asset market trading (Holt, 2006, Chapter 11). We do not expect that many sales offers would be received, but the added quantity uncertainty would probably make it more difficult for bidders to collude in an attempt to drive the price of emission allowances down.

3) Auctions that Combine Different Vintages of Allowances

There is nothing in the structure of two separate uniform-price sealed-bid auctions that prevents the price of the later vintage from exceeding the price of the earlier one. This would be an inefficient outcome and would never be expected to occur in a secondary market for emission allowances.

Combined Vintage Auction

Our report for RGGI recommended two auctions per quarter, one for a fraction of the current vintage and one for a fraction of a future year's vintage. With two auctions occurring simultaneously, it is possible that the price of the later vintage could exceed the price of the earlier one. Arbitrage in spot markets would eliminate a price inversion, but bidders who paid more for a later vintage in an auction would be frustrated. The resulting allocation is likely to be inefficient since allowances are not awarded to those who are presumably willing to pay the most.

Since RGGI allowances are bankable, a bid for a later vintage should be treated as a request to purchase *either* a later vintage *or* an earlier vintage, whichever is less expensive. Interpreting bids in this manner prevents a price inversion in which the uniform price for the later vintage is higher than the price for the earlier vintage with more "convenience value." The four-step process below describes a simple procedure for combined vintage auctions that implements this idea.[1] This procedure can be easily generalized for cases where more than two vintages are auctioned simultaneously.

1) Instead of ranking bids separately for each of the two vintages being sold in a given quarter, all bids for both vintages would be pooled and ranked jointly.
2) Starting at the high end, bids would be awarded the vintage requested if possible. But if the next bid in sequence is for the later vintage and no block is available, then that bidder would receive the earlier vintage if block(s) remain.[2] In contrast, a bid for the earlier vintage would never be awarded a later vintage.
3) This process would continue until all available blocks of each vintage are allocated, and the highest rejected bids for each vintage would then be calculated.
4) The clearing price for the later vintage is the highest rejected bid for that vintage, as for a standard uniform-price auction. The clearing price for the earlier vintage, however, would have to be high enough to "exclude" rejected bidders for either vintage. Therefore, the clearing price for the earlier vintage would be the maximum of the highest rejected bids for either vintage, which ensures that prices will not be inverted.

Advantages:

- The procedure is computationally simple and intuitive: simply transfer supply of allowance blocks from early-vintage bidders to later-vintage bidders as long as the marginal bid for the later vintage is higher. Supply transfers in the other direction would not be permitted, since bidders for an early vintage would not be willing to pay the same bids for a later vintage with lower convenience value.
- The procedure is minimally intrusive and would not affect on allocations if bids for the separate auctions would have yielded a lower price for the later vintage.
- If a price inversion would have occurred in separate auctions, the proposed modification would tend to raise economic efficiency by allocating allowances to those who bid the most, under the reasonable assumption that a bid for a later vintage is really a bid for either vintage, whichever is less expensive.
- The modified procedure does not complicate bidders' decisions; they do not have to bid piecewise-linear functions, discounts, contingent bids, etc.

Numerical Examples: The following table contains two simple examples of auctions with 3 blocks for each of two vintages, 2009 and 2012.

- Example 1 shows how separate auctions may yield an inverted pattern, with prices of $1 for 2009 and $4 for 2012. In the modified auction, a $4 bid for a 2012 block (which would have been rejected in separate auctions) is awarded a

[1] See NYISO (2007) comments on NYSERDA proposed rule 507 for another approach for mitigating this problem that involves specifying premiums for the earlier vintage when placing bids for a later vintage. Then, permits would be allocated to maximize surplus. This would be an added step that would complicate the bidding procedure somewhat, but there could potentially be small efficiency gains over our proposed solution.

[2] A slight modification to this procedure would be to award the earlier vintage to the highest bidder in the later vintage auction rather than the marginal bidder.

2009 block in the row with an asterisk. The highest rejected bids *after* the reallocation determine prices ($3 for 2009, $2 for 2012) that are not inverted.

- Example 2 yields the same reallocation in the modified auction, but the pricing rule has an additional effect in this case. The price for the 2009 block in the modified auction is raised from the highest rejected bid of $2 to the $3 level needed to "exclude" the rejected bid of $3 for the later vintage.

Example 1	**Example 2**
Ranked Bids in Separate Auctions:	**Ranked Bids in Separate Auctions:**
2009: $**6, $4, $3,** rejected: $1, $1	2009: $**6, $4, $2,** rejected: $1, $1
2012: $**5, $4, $4,** rejected: $4, $2	2012: $**5, $4, $4,** rejected: $4, $3
Inverted Prices:	**Inverted Prices:**
2009 price = $1 (highest rejected 2009)	2009 price = $1 (highest rejected 2009)
2012 price = $4 (highest rejected 2012)	2012 price = $4 (highest rejected 2012)
Ranked Bids in Modified Auction:	**Ranked Bids in Modified Auction:**
$6 for 2009, awarded 2009	$6 for 2009, awarded 2009
$5 for 2012, awarded 2012	$5 for 2012, awarded 2012
$4 for 2009, awarded 2009	$4 for 2009, awarded 2009
$4 for 2012, awarded 2012	$4 for 2012, awarded 2012
$4 for 2012, awarded 2012	$4 for 2012, awarded 2012
$4 for 2012, awarded 2009*	$4 for 2012, awarded 2009*
$3 for 2009, highest rejected for 2009	$3 for 2012, highest rejected for 2012
$2 for 2012, highest rejected for 2012	$2 for 2009, highest rejected for 2009
$1 for 2009, rejected	$1 for 2009, rejected
$1 for 2009, rejected	$1 for 2009, rejected
Prices in the Modified Auction:	**Prices in the Modified Auction:**
2009 price = max{2,3} = $3	2009 price = max{2,3} = $3
2012 price = $2 (highest rejected 2012)	2012 price = $3 (highest rejected 2012)

*This is a bid for a 2012 block that is awarded a 2009 block.

References

Cramton, Peter (2007) "Comments on the RGGI Market Design."

Goeree, J. K., T. Offerman, and R. Sloof (2006) "Demand Reduction and Preemptive Bidding in Multi-Unit License Auctions," Discussion Paper, Caltech.

Holt, C. A. (2006) *Markets, Games, and Strategic Behavior*, Boston: Addison-Wesley.

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

NYISO (2007) "Comments of the New York Independent System Operator
On New York State Energy Research and Development Agency Proposed Rule 507 CO2
Allowance Auction Program" December 21, 2007