An Experimental Analysis of Auctioning Emission Allowances Under a Loose Cap

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

The direct sale of emission allowances by auction is an emerging characteristic of cap-and-trade programs. This study is motivated by the observation that all of the major implementations of cap-and-trade regulations for the control of air pollution have started with a generous allocation of allowances relative to recent emissions history, a situation we refer to as a "loose cap." Typically more stringent reductions are achieved in subsequent years of a program. We use an experimental setting to investigate the effects of a loose cap environment on a variety of auction types. We find that all auction formats studied are efficient in allocating emission allowances, but auction revenues tend to be lower relative to competitive benchmarks when the cap is loose. Regardless of whether the cap is tight or loose, the different auction formats tend to yield comparable revenues toward the end of a series of auctions. However, aggressive bidding behavior in initial discriminatory auctions yields higher revenues than in the other auction formats, a difference that disappears as bidders learn to adjust their bids closer to the cut-off that separates winning and losing bids.

Key Words: auction, carbon dioxide, greenhouse gases, allowance trading, Regional Greenhouse Gas Initiative, RGGI, cap and trade

The emerging trend toward the sale of emission allowances as a major method, or even the primary method, for allocating allowances into the economy has focused attention on the implications of various methods for selling allowances.[1] This is especially true given the considerable economic value of the new environmental assets. The U.S. Congressional Budget Office estimates that the aggregate value of allowances under a national greenhouse gas cap-and-trade program would be between $50 billion and $300 billion per year (in 2007 dollars) by 2020. While there have been some experiments with other forms of sale, economic theory, the weight of recent experience, and current policy proposals lead us to expect that auctions will emerge as an important mechanism for selling allowances. It is widely accepted that the performance of different types of auction mechanisms will vary with the institutional context. Taken together, these factors argue strongly for the importance of investigating the performance of different auction designs in the context of an allowance trading program.

This study is motivated by the observation that all of the major implementations of cap-and-trade regulations for the control of air pollution have started with a generous allocation of allowances relative to recent emissions history. At least at the

William Shobe is Director of the Center for Economic and Policy Studies and Charles Holt is Professor in the Department of Economics at the University of Virginia in Charlottesville, Virginia. Dallas Burtraw and Karen Palmer are Senior Fellows at Resources for the Future in Washington, D.C. Jacob Goeree is Professor at the University of Zurich in Switzerland. Erica Myers is a Ph.D. student in the Department of Economics at the University of California in Berkeley, California.

This research was funded, in part, by the New York State Energy Research and Development Authority, the National Science Foundation (Grant No. SES 0098400), the University of Virginia Bankard Fund, and the Simmons Foundation. The authors wish to thank Andrew Barr, A.J. Bostian, Ina Clark, Kendall Fox, Courtney Mallow, Lindsay Osco, and Sara St. Hilaire for research assistance.

This paper was presented as a selected paper at the workshop "The Use of Experimental Methods in Environmental, Natural Resource, and Agricultural Economics," organized by the Northeastern Agricultural and Resource Economics Association (NAREA) in Burlington, Vermont, June 9–10, 2009. The workshop received financial support from the U.S. Environmental Protection Agency, the USDA Economic Research Service, and the Farm Foundation. The views expressed in this paper are the authors' and do not necessarily represent the policies or views of the sponsoring agencies.

[1] The shift toward selling allowances and away from the past practice of allocating them for free to emitters has considerable consequences for both economic efficiency and political support for cap-and-trade programs. See Burtraw et al. (2005).

outset, the supply of allowances has been plentiful relative to the number of allowances needed for compliance. In some cases the supply has remained plentiful, while in others the supply has become progressively tighter. We use the term "loose cap" to refer to cases where the supply of emission allowances is relatively high, that is, relatively close to the unconstrained market equilibrium. This might be relevant in absolute terms, if the emissions constraint does not bind at least for a short period of time, but more generally a loose cap is a relative concept. We investigate the performance of alternative types of auctions when the cap is loose and the expected auction price is low.

The relative slackness of an emissions cap could have implications for the preferred design of an allowance auction. Smith (1967) found that, in an experimental setting where bidders in a sealed-bid auction shared a common, but imperfectly known, value for multiple units of a good, a loose cap produced a divergence in the performance of different auction types. Specifically, Smith found that uniform-price auctions generated more revenue than discriminatory-price auctions when excess demand was low, but that uniform and discriminatory auctions generated comparable levels of revenue when excess demand was high.

In this paper, we investigate the effects of a loose cap environment on a variety of auction types. We compare single-round, sealed-bid auctions, both discriminatory and uniform price, with the multi-round, ascending-price, English clock auction. In single-round, sealed-bid auctions, bidders submit a demand schedule for blocks of allowances so that bids have both price and quantity dimensions. In multi-round auctions, bidders submit quantities at a given price. If demand exceeds supply, the price is raised (like the advancing of a clock) until demand is equal to supply. We investigate the performance of the English clock auction both with and without excess demand information provided to bidders between rounds.

We find that whether the cap is tight or slack, all auction formats studied are nearly 100 percent efficient in allocating emission allowances, with no systematic differences among them—e.g., allowances are distributed initially in the auction to those who value them most. As we would expect in the less competitive environment of a loose cap, the amount of revenue raised in the auction appears to be less, relative to a competitive (Walrasian) prediction determined by the supply and demand for permits, than revenue under a tight cap for all auction types. In terms of revenue generation, differences between auction formats are relatively minor. However, one notable effect of the auction design is that aggressive bidding in discriminatory auctions yields higher revenue early in a series of auctions, a pattern that tends to diminish later as bidders learn more about the likely level of the highest rejected bid.

Background and Literature

Several cap-and-trade programs for controlling air emissions have been characterized by loose caps, particularly during their early years of operation (Ellerman et al. 2000). More generally, research by Harrington, Morgenstern, and Nelson (2000) suggests that environmental regulations often end up being less costly than was anticipated at the time of their adoption and that overestimation of baseline emission levels and failure to take account of the emission-reducing effects of technological change are important factors that contribute to that overestimation.

Probably the most famous case of a loose cap is Phase I of the European Union Emissions Trading Scheme (EU ETS). This program was implemented prior to the availability of an accurate inventory of CO_2 emissions in the EU. Once it became clear that the number of allowances allocated was more than needed for compliance, the price of EU ETS allowances crashed. The fall was exacerbated by the limited lifespan of Phase I allowances, which were not bankable into Phase II. One auction of CO_2 allowances under Phase I closed at a price of 0.7 euros per ton of CO_2, a fraction of the price just weeks earlier (Ellerman and Buchner 2007).

The first incentive-based program to be implemented anywhere in the world on a large scale was the phase-down of lead in gasoline in the United States in the mid-1980s. The regulations permitted trading and banking of lead permits through inter-refinery averaging of the lead content of a gallon of gasoline. The standard started at 1.1 grams of lead per gallon in 1982, was lowered to 0.5 in 1985, and finally was reduced to 0.1 in 1986 (Kerr and Newell 2003).

In California, the Regional Clean Air Incentives Market (RECLAIM) program to cap emissions of

NO_x in the Los Angeles Basin began in 1994. Under this program, regulated facilities were free to pick their baseline emissions level from a time period of strong economic activity, and as a result the aggregate emissions cap was not binding during the early years of the program (Burtraw et al. 2005). The Emissions Reduction Market System for volatile organic materials in Chicago has been persistently slack, with emissions falling far below the cap between 2001 and 2005 (Evans and Kruger 2007). A number of factors have contributed to the slackness of the Emissions Reduction Market System, including inflated baseline emissions, closure of regulated facilities, and interactions with regulations directed at other air pollutants.

The U.S. sulfur dioxide emissions trading program initiated in 1995 was specifically designed to be relatively slack at first, with a significant tightening of the cap in a second phase. Changes in technology, fuel markets, and regulations of the railroads contributed to a decline in compliance cost by the first year of implementation that was not anticipated when the program was enacted in 1990 (Burtraw 2000, Carlson et al. 2000). Only after trading started in 1993 in advance of implementation did it become clear just how slack the initial cap was. An annual revenue-neutral auction that was first held in advance of the compliance period played an important role in the discovery of compliance costs and the associated allowance price that ultimately obtained in the market. Low prices, thin trading, and large amounts of allowance banking characterized the early market for SO_2 allowances (Ellerman et al. 2000).

In January of 2009, ten northeastern states from Maryland to Maine launched the Regional Greenhouse Gas Initiative (RGGI), the second mandatory emissions cap-and-trade program for greenhouse gas emissions in the world and the first in the United States. For the years 2009 through 2014, the RGGI program caps total annual CO_2 emissions from electricity generators in the region at a level approximately 4 percent above the average annual CO_2 emission levels in the years 2000–2004. Beginning in 2015, the cap will decline about 2.5 percent per year for four years, until it is 10 percent below the initial cap by 2019. In addition to allowing full trading of allowances, the program provides flexibility by allowing for 3-year compliance periods, with the first period extending from 2009 through 2011, full allowance banking into future compliance periods, and limited use of emission offsets.

The RGGI cap on carbon emissions was set with the idea that it would be "loose" in the early years of the program, so as not to stifle economic growth or cause shortages in the electricity market. However, as the first CO_2 cap-and-trade program in the United States, the RGGI program is subject to much uncertainty. One of the many sources of uncertainty is the stringency of the annual CO_2 emissions cap compared to actual emission levels in the absence of a cap. Several recent analyses (Environment Northeast 2008, New Carbon Finance 2007, Point Carbon Research 2007) have compared the trend in CO_2 emission levels and projected emission levels from electricity generators in the region to the level of the emissions cap. The evidence suggests that, at the outset of the program, projected annual emissions could fall short of the initial annual RGGI emissions cap of 188 million tons.

Even before the 2008–2009 recession, predictions suggested that the RGGI caps might not be binding in the initial years of the program. Predictions of slack demand were based on the growth in natural gas fired generation in the region,[2] resulting both from greater use of natural gas in dual-fueled boilers that can burn both oil and natural gas and, importantly, the increase in new, highly efficient, natural gas fired capacity in the region. The relative costs of different fossil fuels helps determine the demand for emission allowances through their effect on the dispatch of different types of generators in the electricity sector. For example, in 2004 a spike in natural gas prices led to greater use of coal to produce electricity, which contributed along with the introduction of the Clean Air Interstate Rule to a marked increase in SO_2 allowance prices (Burtraw et al. 2005). Two years later, total emissions of SO_2 from regulated generators fell by over 800,000 tons relative to emissions in 2005. Nearly 30 percent of that reduction (or approximately 3 percent of the total annual supply of allowances) was attributable to the substitution of natural gas (which emits no SO_2) for oil (which does) at dual-fueled generating units in response to increases in the relative

[2] Weather has also been mentioned as a contributing factor. It is not yet known whether this is a transitory or longer-term phenomenon.

price of fuel oil (U.S. Environmental Protection Agency 2007). An increase in demand for CO_2 emission allowances of similar magnitude could result if a major nuclear plant in the region were to unexpectedly go offline for a period of time.

Another factor that will influence demand for emission allowances is the level of overall demand for energy. If a cap-and-trade program is coupled with programs to enhance energy efficiency, as envisioned in the CO_2 programs in RGGI and in California (California Air Resources Board 2008), the level of demand for allowances will depend importantly on the effectiveness of those energy efficiency programs and policies. In RGGI, the memorandum of understanding among the ten participating states specifies that at least 25 percent of the value of emission allowances must be directed toward strategic energy investments such as renewables and energy efficiency, and most of the states plan to use at least this portion of the RGGI allowance auction revenue to fund energy efficiency programs. Simulation modeling suggests that increasing the share of emissions-allowance revenue dedicated to funding efficiency programs could reduce RGGI allowance prices by as much as 9 percent relative to a baseline of only 25 percent of RGGI revenue being dedicated to energy efficiency (Ruth et al. 2008).

In the case of RGGI, there are factors that will help to mitigate the type of allowance-price collapse that happened in the European Union allowance market at the end of the first commitment period. First, there is a reserve price in the RGGI allowance auctions of $1.86 per ton, which means that the RGGI states will sell no allowances in the auction at a price below $1.86. Second, and unlike Phase I of the EU ETS, the RGGI program allows full banking of allowances and the possibility of purchasing allowances at a low price today and saving them for use when economic growth leads to greater electricity demand. In fact, in the first four RGGI auctions, all of which occurred during a serious recession, allowance prices closed well above the reserve price of $1.86, at $3.07, $3.38, $3.51, and 3.23, respectively.[3] Even as demand faltered and prices fell due to low utility emissions during 2009, the reserve price was never binding. The ability to bank allowances into future, post-recession years helps keep the price of allowances above the reserve price even in years of relatively slack demand.

The relevance of this slack demand to auction design arises from a developing expectation that, in the future, most cap-and-trade programs, and in particular those for controlling CO_2 emissions, will provide for the auction of a significant share of allowances. Unlike what has occurred in previous emission cap-and-trade programs, the states participating in the RGGI program are, collectively, planning to auction roughly 90 percent of the emission allowances created by the program.

The RGGI states decided to expand the role for auctions, in part because they recognized that giving allowances away for free would not prevent electricity producers in the region from folding the value of emission allowances into their bids in the power market. The region has moved almost entirely away from cost-of-service pricing to market-based pricing in electricity markets, which means that the opportunity costs of emission allowances would be passed on in electricity prices even if allowances were received for free by the regulated firms.

In contrast, most prior emissions trading programs, including the SO_2 and NO_x trading programs in the United States and the first two phases of the European Union's Emissions Trading Scheme, began by granting the emission allowances to current emitters for free. Recent regulatory and legislative proposals indicate that the RGGI cap-and-trade program presages a trend of increasing reliance on auctions for allowance allocation. For example, current proposals in the EU provide for selling a substantial majority of the allowances used by electricity generators in the third phase of its program beginning in 2013. At the time of this writing all major proposals for a U.S. federal greenhouse gas cap-and-trade program envision a significant role for auctions, although some prominent proposals provide, at least initially, for a split between auctions and free allocation.

As allowance auctions become more prevalent, their performance under different market settings takes on increasing importance. Given the possibility of low excess demand especially at the outset of allowance markets, the choice of auction

[3] For the auctions from March 2009 forward, an auction was also held for a later vintage of allowance, which cannot be used prior to 2012. These allowances sold for lower prices, reflecting their lower "convenience value."

design should include consideration of the impact of low excess demand on auction performance.

Procedures

In the RGGI program, emission allowances are distinguished by the first date at which they can be used for compliance, and they can be used at any time after that. Allowances with the same vintage are identical, and therefore auctions for emission permits are multi-unit auctions. In the experiment, the term "permit" is analogous to allowances in RGGI, each bidder is assigned several units of production capacity, and each unit of production from a particular type of capacity unit requires some number of emission permits, which varies among participants to reflect a distribution of combustion technologies. Bidders' values for permits are determined by the profit margins on their production capacity and by the number of allowances needed to cover the production activity.

The auction formats that we consider are distinguished in two dimensions: by whether or not there are multiple rounds of bidding, and by whether or not all bidders purchase permits at the same "uniform" price. In the single-round auctions, bidders submit sealed bids for blocks of permits, and it is possible to make different bids for each block. These bids are then ranked from high to low as a pseudo demand function, and then the vertical supply of permits being auctioned determines which bids are accepted. In a "discriminatory price" auction, successful bidders pay their own bid prices, whereas in a single-round "uniform-price" auction, all pay the same market-clearing price, which is set at the highest rejected bid. In both formats, bids are required to exceed a pre-announced reserve price. The third format we consider is a multi-round "clock auction" in which bidders state how many blocks they wish to purchase at an announced price, and in which in each round the price is raised sequentially, as if by hands of a clock. The clock price begins at the reserve price and stops rising when aggregate bidder quantity demanded is less than or equal to the auction quantity supplied. In each round, bidders can reduce their quantity as the price rises, or leave demand unchanged, but they cannot increase the quantity of their bid according to the auction "activity rules" that force bidders to participate actively from the beginning. To maintain comparability with the other auction formats, prices with the clock are incremented by the same fixed increment that corresponded to possible bid increments in the single-round formats. The experiments implemented two versions of the clock: one in which bidders were not told the aggregate demand quantity after each round, and one in which they were provided with this information.[4]

Burtraw et al. (2009) report an experiment with these auction formats under treatment conditions that were designed to facilitate tacit or explicit collusion, e.g., the presence of a chat room or the existence of a secondary market where trades can be arranged to adjust for failure to obtain allowances in a collusive phase. In all of the auctions reported in that paper, however, the cap on emissions was tight in the sense that the available supply of permits in each auction was only about two-thirds of what would be demanded at a price of zero. In this paper, we report two series of auctions: a baseline series where supply is two-thirds of the quantity that would be demanded at a zero price, and a loose cap series with a cap that represents about 90 percent of the quantity that would be demanded at a zero price.

In a loose cap environment, most bids will end up being accepted, and the resulting reduction in competition can serve as a useful "stress test" of auction formats. In a uniform-price auction, there may even be a role for the exercise of unilateral market power if the cap is so loose that a single bidder can profit from bidding lower on allowances for marginally profitable capacity units in the hopes of lowering the clearing price on a large number of other, more profitable units. This strategy is referred to in the literature as "demand reduction." In a clock auction with low excess demand to begin with, a small amount of demand reduction by one or more bidders will stop the clock at low prices. Similarly, there is less risk to bidding low on marginal units in a multi-unit discriminatory auction with a loose cap since a higher proportion of bids will be accepted, and the resulting bid reductions may cause revenue to be lower than in uniform-price auctions. This revenue comparison might be reversed if there is

[4] The clock auction implemented in the recent Virginia NO_x auction did not provide excess demand information (Porter et al. 2009), and the auction administrators felt that this decision may have prevented the clock from stopping before it did (Holt et al. 2007).

demand-withholding in uniform-price or clock auctions.

The experimental design implemented an asymmetry motivated by the difference between high emitters (e.g., those using coal) and low emitters (those using natural gas). Each laboratory auction involved 6 bidders, with 3 "high users" requiring 2 permits per unit of production capacity used, and 3 "low users" who were required to remit 1 permit per unit of production capacity used. All bidders were given 5 units of capacity, regardless of whether they were low or high users, reflecting the approximate equal proportions of electricity production by each type of producer. Thus, low users would demand 15 permits (= $3 \times 5 \times 1$) and high users would demand 30 permits (= $3 \times 5 \times 2$), for a total demand of 45 permits at a price of zero. For the loose cap treatment, 40 permits were sold in each auction, and for the tight cap treatment, 30 permits were sold in each auction. Each unit of capacity could be used to produce a unit of a product that sold for a known price. Since permits could not be banked in this experiment, their values were determined by the difference between the product price and the production cost, divided by the number of permits required. Bidders received randomly determined production costs associated with each unit of capacity; these costs were draws from uniform distributions—[2, 6] for high users, and [5, 10] for low users—with new draws made for each successive auction. The permit values resulting from a given set of cost draws can be arrayed from high to low like a demand curve, as shown in Figure 1, where the auction supply is represented as a vertical line at the supply of 40 permits (loose cap).

In Figure 1, the intersection of the value array and the auction supply determines a Walrasian price, which will vary from one auction to the next depending on the actual cost draws. Laboratory sessions involved 6 bidders who participated in a series of 8 auctions over the course of an hour (including the reading of instructions). Both loose and tight cap treatments involved uniform, discriminatory, and clock auction formats. In addition, the loose cap treatment included a clock auction with full information about excess demand at the end of each round of bidding. Each combination of auction format and cap (loose or tight) involved 6 experimental sessions, each with 6 bidders. The random cost draws were balanced across treatments, by using a given random number "seed" for the first session of each auction type, a second random number seed for the second session of each auction type, etc. There were 252 subjects who participated in the study. All were students at the University of Virginia, and were paid their earnings in cash at the end of each session. Participants received a $6 initial payment, and earnings in the auctions ranged from $15 to $45 after being converted (at a rate of 30 cents per experiment dollar). The auctions were run using a web-based program that is publicly available on the Auctions Menu of the Veconlab site.[5]

For each auction, we determined the maximum economic surplus that could be achieved given the fixed product price, the individual cost draws, and the number of permits being auctioned. This surplus is the area below the value array and to the left of the auction supply. In an actual auction, variations in bidding behavior may cause some permits to be purchased by bidders with low values, and the actual efficiency is determined by dividing the value of the resulting allocation by the maximum possible value. A second measure of auction performance is the amount of revenue that it raises. If all bidders bid their full values in a discriminatory (pay-as-bid) auction, the seller would receive the maximum revenue amount, which is the area below demand and to the left of the auction quantity, i.e., the entire economic surplus. We will use a more realistic benchmark for revenue: the amount of money that would be raised if all permits sold at the market-clearing "Walrasian" price determined by the intersection of the supply and demand curves as depicted in Figure 1 for the case of a loose cap.

Results

Figure 2 shows the average prices for each of the 8 auctions with a tight cap. Each solid line corresponds to a different auction format, and the dashed line shows the Walrasian price predictions. All prices represent averages across the 6 sessions. All these prices are close to the Walrasian prediction, and there is no obvious trend in price levels. The small price differences do not result in significant differences in auction revenues as determined by a Wilcoxon matched-pairs

[5] At http://veconlab.econ.virginia.edu/admin.php.

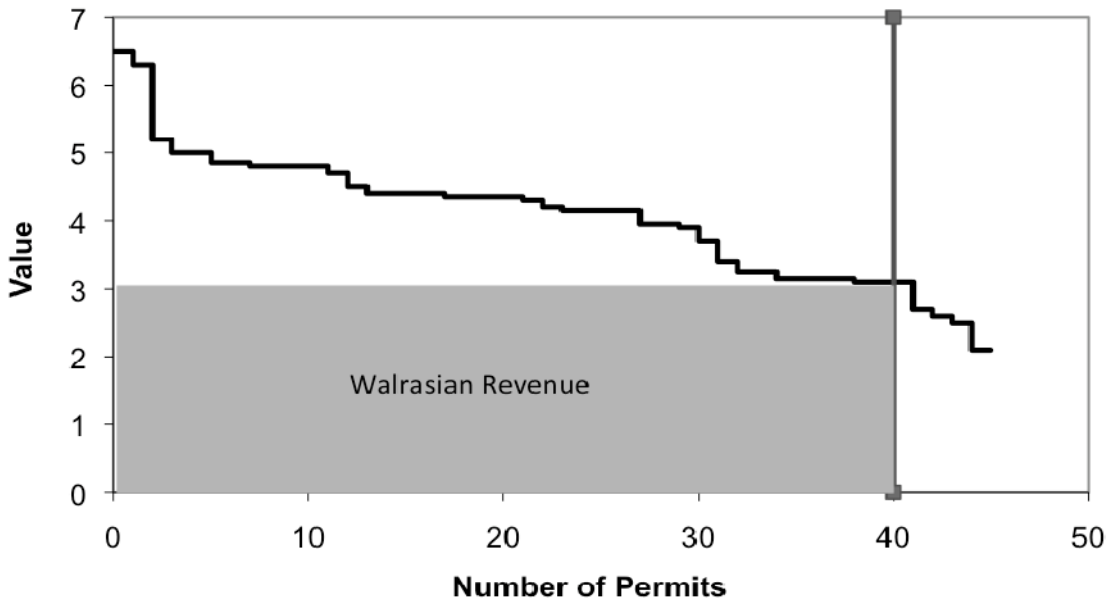

Figure 1. Permit Values and Walrasian Equilibrium with a Loose Cap

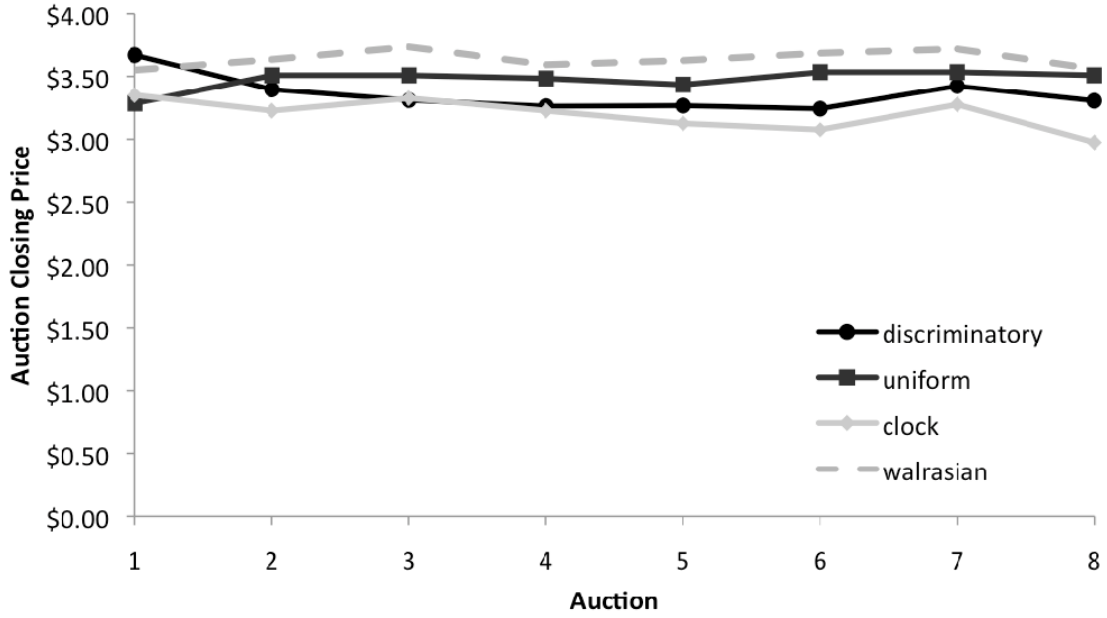

Figure 2. Average Price Paid per Permit: Tight Cap

test using revenue averages for each session.[6] To summarize, differences in auction formats do not have much impact on revenues in the competitive, tight cap environment.

Figure 3 shows the average price for each of the auction formats in the loose cap environment. This dynamic representation reveals an interest-

[6] Using data from the final four auctions, the p-values for pairwise comparisons between auction formats are all greater than 0.10.

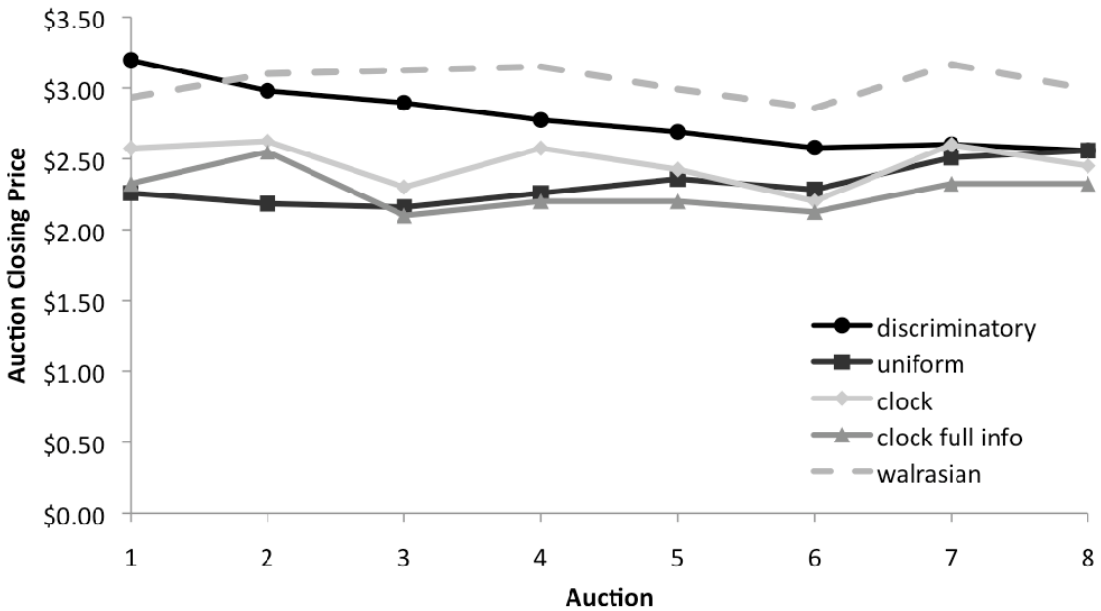

Figure 3. Average Price Paid per Permit: Loose Cap

ing pattern in which the discriminatory auction yields higher average prices especially in early auctions, but this difference appears to taper off by the end of the series. Using data from the first four auctions in each session, the discriminatory format raises significantly more revenue than the other three formats.[7] In the last four auctions, the only significant difference is that the discriminatory auction raises more revenue than the clock auction with full information (W= 0, p= 0.031). There is no difference between the revenues raised under the two types of clock auctions in either case (W= 6, p= 0.438, for the first four auctions, and W= 5, p= 0.313, for the last four auctions).

Table 1 summarizes the results for tests of differences between auction types within a given "cap tightness" regime. Given the null hypothesis of no differences between auctions under a given regime, we cannot reject the null hypothesis for any auction comparisons under the tight cap. That is, evidence does not suggest a difference between auction types under a tight cap. Under the loose cap, only the comparison of discriminatory sealed-bid to the full-information clock gives a statistically significant difference; we reject the null hypothesis with p= 0.031, and conclude that discriminatory auctions raise more revenue than the full-information clock.

Table 1. Differences in Auction Revenues Within a Regime

Regime	Test of H_0 that auction types raise the same revenue (last 4 observations)
Tight	Cannot reject H_0 for any pairwise comparison of auction types.
Loose	Discriminatory auction raises more revenue than the full-information clock auction (reject H_0 with p=0.031). Cannot reject H_0 for any other pairwise comparison.

Some interesting properties of the sealed-bid auctions are revealed by an analysis of individual bid data and the way that bid patterns evolve over time. Figure 4 shows the bid data for the six uniform-price sessions. Bid values are on the vertical axis, and permit values are on the horizontal axis. Individual bids are shown for the first four auctions (left) and the last four auctions (right). Larger data points represent higher frequencies of that bid/value combination. If bidders are bidding their true values for permits, the data would fall

[7] In the first four auctions, the signed rank test statistic comparing the discriminatory-price auction to each of the other auction formats considered yielded a test statistic of 0 and a p-value of 0.031 in each case.

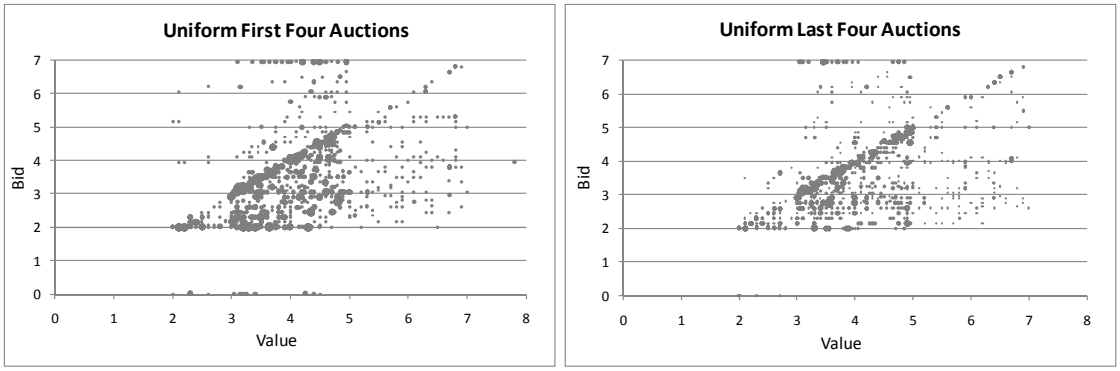

Figure 4. Uniform Auction Individual Bids

along the 45 degree line through the origin. In both panels, there is a clear density of bids along the 45 degree line, showing that in uniform-price auctions, many bidders bid their true values. Also notable is the instance of many bids above the 45 degree line, particularly in the mid-value range, where bids are less likely to be marginal. There are few differences between the early and late bid distributions, indicating that bidding behavior is fairly consistent throughout the experimental sessions. In the final half, however, there is a slight increase in the density of bids at value, and a slight decline in bids at the reserve price for these uniform-price auctions.

Figure 5 shows the bid data for the six discriminatory sessions for both the first four auctions (left) and the last four auctions (right). It is clear from comparing these figures with the comparable panels in Figure (4) that bidders implement different strategies in the two different auction formats. In the discriminatory auctions, there are far fewer bids on the 45 degree line, and the bids are more uniformly distributed between $2 and $3, the typical clearing price. In discriminatory, pay-as-bid auctions, bidders attempt to bid just above the market-clearing price for their higher-valued units in order to maximize the amount of surplus that they capture. The differences in bidding strategy between the two auction formats may account for the dynamic trend that we see in the loose cap data, where the discriminatory auction has higher average prices than the other auction types initially, but where over time the prices drop with successive auction rounds. It may take several auctions for bidders to "discover" the clearing price and bring down the average price paid per permit. This can be seen by comparing the bid scatter graph for the first 4 auctions with that of the final 4 auctions, as shown in Figure 5. In the early auctions, bids are more widely distributed between the reserve price and the bidders' values, but in the later sessions, as bidders learn more about the likely cut-off price, the bids tend to be more clustered in the $2–$3 range.

Figure 6 depicts the proportion of the Walrasian revenue that was raised in the final four auctions of each session, with sessions grouped by auction format and tight cap versus loose cap environment. Similarly, the bars in Figure 7 show the efficiency proportions for the final four auctions in each session in both the loose cap and tight cap settings. Each bar represents a session, where the first session listed in each group was done with the first random number seed, etc. There are no clear differences among treatments or auction formats in terms of efficiency, with all values close to 100 percent. In contrast, revenue percentages all appear well below the Walrasian benchmark of 1, which was the expected result given the non-competitive nature of the experimental setting.[8] For the pooled data across auction types, the two-sample Wilcoxon rank sum test ($W = 265$ and $p = 0.031$) allows us to reject, with better than 95 percent confidence, the null hypothesis that revenues are the same under a

[8] In a series of experiments that we ran under more competitive conditions with 12 bidders and a tighter cap (equal to 70 percent of demand for permits at a zero price), we found revenues to be very close to Walrasian levels for uniform, discriminatory, and clock auction formats, with no systematic differences across auction types (Holt et al. 2007).

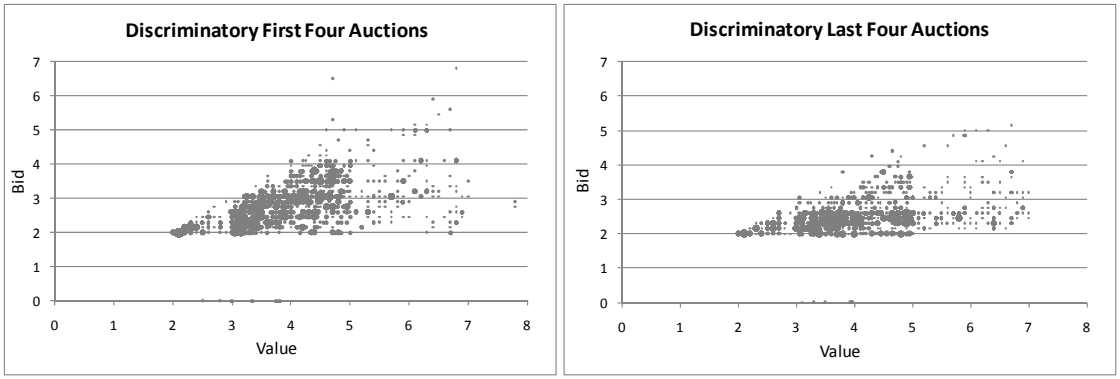

Figure 5. Discriminatory Auction Individual Bids

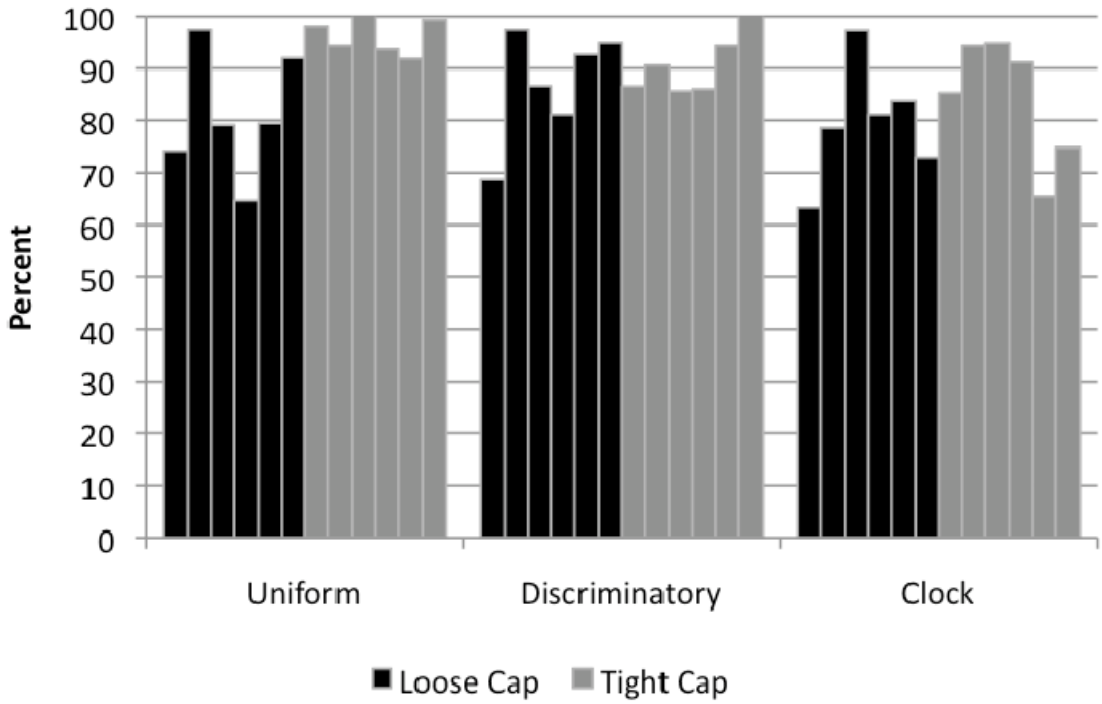

Figure 6. Auction Revenue as a Proportion of Walrasian Revenue

loose cap as under a tight cap. As expected, a tight cap results in more revenues than a loose cap.

This effect is not consistent across auction types. Using the Wilcoxon rank sum test for differences in the proportion of Walrasian revenue raised between the tight and loose cap within each auction format, differences arise. Each session was an observation where the proportion of Walrasian revenue averaged across the last 4 auctions served as the measure for comparison. We are unable to reject the null hypothesis of no difference in the proportion of Walrasian revenue raised between the tight and loose cap environment for the discriminatory ($W = 38$, $p = 0.937$) and clock ($W = 34$, $p = 0.485$) auctions. However, we found that the uniform-price auction raised a larger proportion of Walrasian revenue in the

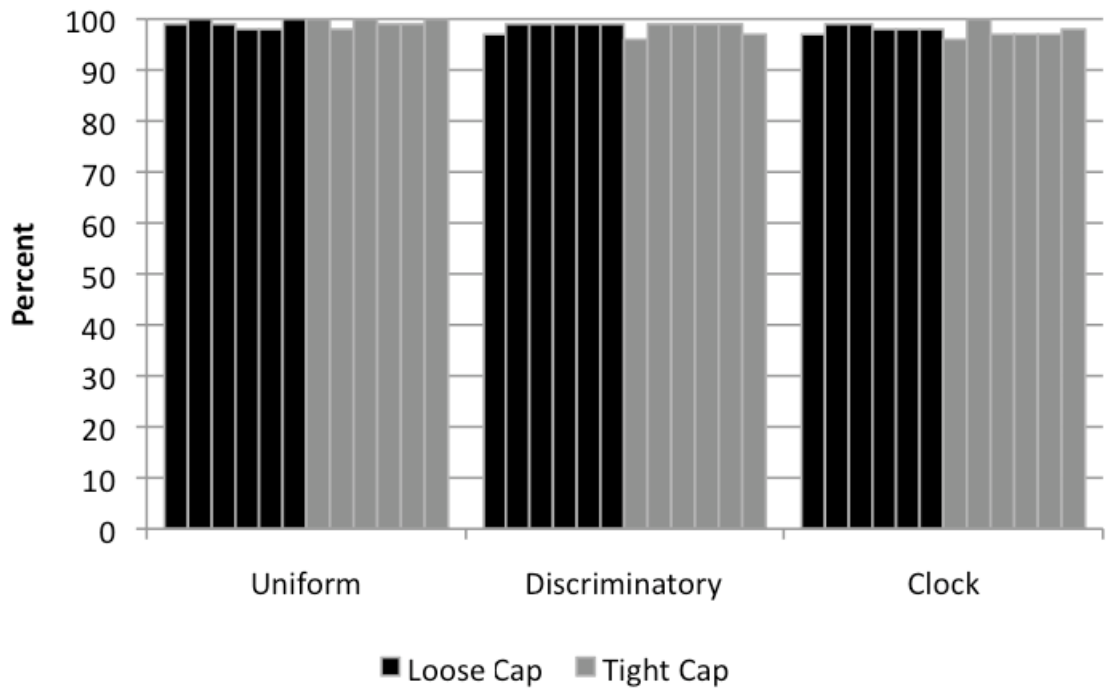

Figure 7. Efficiency

tight cap environment relative to the loose cap environment (W= 24, p= 0.015). The results are summarized in Table 2.

It is well understood that a bidder in a uniform-price auction may have incentive to shade bids on marginal units in an effort to lower the clearing price on the bidder's infra-marginal bids. This depends on the bidder's assessment of the probability that a reduction in a bid will lower the clearing price and by how much. The gain from such a reduction would depend on the number of winning bids (with bids close to the expected closing price) that the bidder will have—the more the accepted bids, the greater will be the value of any price reduction resulting from shading bids on marginal units. Earlier studies that have investigated the exercise of market power explicitly controlled market power and made the extent of this power known to subjects by providing detailed information about the market structure. In our experiments, the extent of market power depends on random elements and is difficult for the bidders to evaluate.

Table 2. Differences in Auction Performance Between Loose and Tight Caps

Auction Type	Result
Pooled data	Revenues greater with tight cap (p=0.031)
Discriminatory sealed-bid	No significant difference
Uniform sealed-bid	Revenues greater under tight cap (p=0.015)
Clock	No significant difference

If the bidders were acting on a correct assessment of their market power, then we would expect that any such bid shading would be positively associated with the size of the expected price movement and the expected number of winning bids. It seems reasonable to expect, then, that we would see bids falling further below their actual value for those allowances close to the closing price, and that users needing more allowances would tend to shade more than those needing fewer allowances.

Figure 8 shows the bid-to-value ratios broken into allowance value ranges. Bids are separated

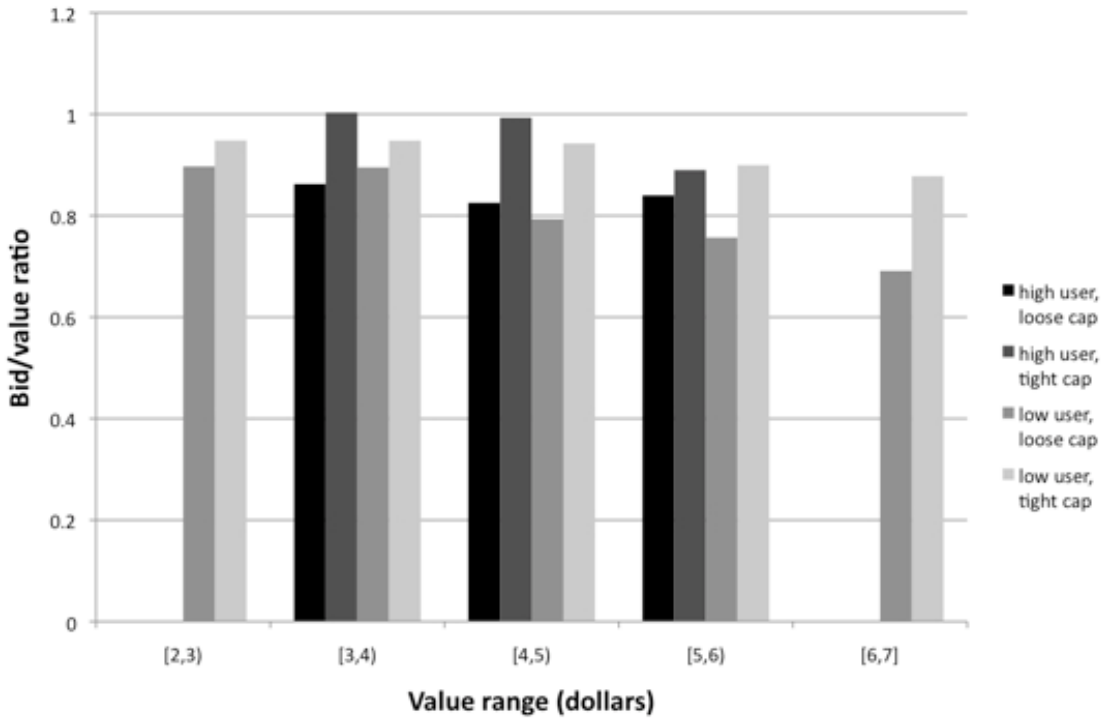

Figure 8. Bid-to-Value Ratios for Uniform Price Auction

by tightness regime and by users with high and low need for allowances in production. It is consistent with the demand reduction hypothesis that, in each value range, the average bid-to-value ratio is lower with a loose cap than with a tight cap. This is true for both low users and high users. However, under the demand reduction hypothesis just described, we might also expect to see greater bid reductions for allowances with values in the $3 to $4 range, since these are the marginally profitable allowances, and this effect should be greater for users needing high numbers of allowances than for those needing low numbers of allowances because there are more infra-marginal units on which to benefit from a reduced price. This second expectation about bid shading is not evident in Figure 8. There does not appear to be a significantly greater difference in bid-to-value ratios in the $3 to $4 range than in other value ranges. Exploring the bidding behaviors that lead to the lower prices for uniform-price auctions under a loose cap would seem to be a fruitful area for future research.

The results reported here contrast with those reported in Smith (1967). Smith found that seller revenues were higher in a uniform-price auction than in discriminatory auctions in settings with moderate numbers of rejected bids, and that this difference is not apparent with high numbers of rejected bids. While these were multi-unit auctions, they are not directly comparable to our results since bidders were allowed to submit only two bids; our design allowed bidders to submit any number of bids for blocks of allowances. Moreover, in Smith's experiments, about a third (8 of 26) of the bids would be rejected even in the treatment with the fewest number of rejected bids—a considerably more competitive environment than the "loose cap" treatment tested in our experiments. As we have already argued, our less competitive set-up is more consistent with what has been observed in the early stages of emission trading regimes "in the wild." Furthermore, the Smith experiment was done in an environment that was motivated by the Treasury bill auctions, in which the prize values to bidders were identical for all units and were randomly determined, i.e., a random common value. The experiments

reported in this paper implement a private value setting to emphasize the private nature of information about pollution control costs.

Conclusion

The initial allocation of emission allowances in a cap-and-trade program is one of the most contentious issues surrounding this form of environmental regulation. Increasingly, policymakers are turning to the sale of emission allowances as the desired approach for distributing some or all of the allowances created by a cap-and-trade program, and to auctions as the best approach for structuring that initial sale. Given the substantial value of emission allowances associated with future cap-and-trade programs, particularly those designed to reduce emissions of greenhouse gases, finding the most efficient auction method and a way to capture the full market value of that newly created public asset is of keen interest to government officials. In addition, federal legislation to establish cap-and-trade programs for CO_2 in the United States and the current proposal for phase III of the EU ETS envision using a portion of the revenue from an allowance auction to promote carbon-free technologies and energy efficiency in an effort to ease the transition to a world with lower CO_2 emissions. In order to be able to fund these efforts to the greatest extent possible, policymakers are keenly interested in having an auction design that will lessen opportunities for collusion and allow allowances to be sold at their full market value.

Past research suggests that the amount of revenue raised in different auction formats could depend on the amount of excess demand in the allowance market; and, given the prominence of relatively loose caps in the early years of past and existing cap-and-trade programs, it is important to test the efficacy of different auction types. We use laboratory experiments to test the performance of four auction types in the presence of a loose cap, including uniform single-round, discriminatory single-round, clock, and clock with excess demand information.

We find that, with a loose cap, all of the auction formats investigated yield a nearly 100 percent efficient allocation of permits. As expected, for all auction types the amount of revenue raised in the auction appears to be less than revenue raised with a tight cap. Given a loose cap, auction type does affect revenues. Discriminatory auctions yielded higher revenues than sealed-bid uniform-price or the clock auction designs. However, the differences among the different auction types eroded over time during a particular auction session. The convergence of results is consistent with learning and price discovery affecting the bids of participants in discriminatory auctions. Bidders in discriminatory-price auctions appear to be more successful in later rounds at pegging their bids to the expected market price. This suggests that the advantage of the discriminatory-price auction would likely not persist in a regime of regular greenhouse gas auctions.

Our experiments did not include an opportunity to bank allowances for use in future periods. The opportunity to bank allowances would support the price of allowances and attenuate the impact of the loose cap. The extent to which banking would alter the performance of auctions under a loose cap may be a fruitful area for future research.

References

Burtraw, D. 2000. "Innovation Under the Tradable Sulfur Dioxide Emission Permits Programme in the U.S. Electricity Sector." Resources for the Future Discussion Paper No. 00-38, Resources for the Future, Washington, D.C.

Burtraw, D., D.A. Evans, A. Krupnick, K. Palmer, and R. Toth. 2005. "Economics of Pollution Trading for SO_2 and NO_x." *Annual Review of Environment and Resources* 30: 253–289.

Burtraw, D., J. Goeree, C. Holt, E. Myers, K. Palmer, and W. Shobe. 2009. "Collusion in Auctions for Emissions Permits: An Experimental Analysis." *Journal of Policy Analysis and Management* 28(4): 672–691.

California Air Resources Board. 2008. "Climate Change Proposed Scoping Plan: A Framework for Change." Available at http://www.arb.ca.gov/cc/scopingplan/document/psp.pdf (accessed March 29, 2010).

Carlson, C., D. Burtraw, M. Cropper, and K. Palmer. 2000. "SO_2 Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy* 108(6): 1292–1326.

Ellerman, A.D., and B.K. Buchner. 2007. "The European Union Emissions Trading Scheme: Origins, Allocation, and Early Results." *Review of Environmental Economics and Policy* 1(1): 66–87.

Ellerman, A.D., P.L. Joskow, R. Schmalensee, E.M. Bailey, and P.K. Monteiro. 2000. *Markets for Clean Air: The U.S. Acid Rain Program*. New York: Cambridge University Press.

Environment Northeast. 2008. "Emissions Trends and the Inaugural Allowance Auction: Regional Greenhouse Gas Initiative." Available online at http://www.env-ne.org/public/resources/pdf/ENE_RGGI_Emissions_and_Auction_Rpt_090308_final.pdf (accessed September 16, 2008).

Evans, D., and J. Kruger. 2007. "Where's the Sky's Limit? Lessons from Chicago's Cap and Trade Program." *Environment* 49(2): 18–32.

Harrington, W., R.D. Morgenstern, and P. Nelson. 2000. "On the Accuracy of Regulatory Cost Estimates." *Journal of Policy Analysis and Management* 19(2): 297–322.

Holt, C., W. Shobe, D. Burtraw, K. Palmer, and J.K. Goeree. 2007. "Auction Design for Selling CO_2 Emission Allowances under the Regional Greenhouse Gas Initiative." RGGI Report, New York State Energy Research and Development Authority, Albany, NY.

Kerr, S., and R.G. Newell. 2003. "Policy-Induced Technology Adoption: Evidence from the U.S. Lead Phasedown." *Journal of Industrial Economics* 51(3): 317–343.

New Carbon Finance. 2007. "Regional Greenhouse Gas Allowances: Going, Going, Gone?" *North America Research Note* (November 2007): 1–13.

Point Carbon Research. 2007. "Emissions Trading in the U.S.: Is RGGI Over-Allocated?" Carbon Market Analyst. Available at http://www.pointcarbon.com/research/cmana/cmana/1.260304 (accessed March 29, 2010).

Porter, D., S. Rassenti, W. Shobe, V. Smith, and A. Winn. 2009. "The Design, Testing, and Implementation of Virginia's NO_x Allowance Auction." *Journal of Economic Behavior and Organization* 69(2): 190–200.

Ruth, M., K. Ross, N. Hultman, J. Mauer, I.D. Valencia, N. Herrmann, K. Palmer, A. Paul, E. Myers, B. Hobbs, D. Irani, J. Michael, and Y. Chen. 2008. "The Role of Energy Efficiency Spending in Maryland's Implementation of the Regional Greenhouse Gas Initiative." A study commissioned by the Maryland Department of the Environment and conducted by the Center for Integrative Environmental Research at the University of Maryland. Available at http://www.cier.umd.edu/RGGI/CIER_RGGI_Energy_Efficiency_Spending_Study%5B1%5D.pdf (accessed March 29, 2010).

Smith, V. 1967. "Experimental Studies of Discrimination versus Competition in Sealed-Bid Auction Markets." *Journal of Business Research* 40(1): 56–84.

U.S. Environmental Protection Agency. 2007. "Acid Rain and Related Programs: 2006 Progress Report." Report No. EPA-430-R-07-011, Clean Air Markets Division, Office of Air and Radiation, U.S. Environmental Protection Agency, Washington, D.C.