DISCUSSION PAPER

Price and Quantity "Collars" for Stabilizing Emissions Allowance Prices

An Experimental Analysis of the EU ETS Market Stability Reserve

Charles A. Holt and William Shobe

RESOURCES
FOR THE FUTURE

Price and Quantity "Collars" for Stabilizing Emissions Allowance Prices: An Experimental Analysis of the EU ETS Market Stability Reserve

Charles A. Holt and William Shobe

Abstract

This paper reports the results of a laboratory experiment with financially motivated participants that is used to compare alternative proposals for managing the time path of emissions allowance prices in the face of random firm-specific and market-level structural shocks. In this setting, market performance measures such as social surplus are enhanced by the use of a price collar (auction reserve price and soft price cap). Comparable performance enhancements are not observed with the implementation of a quantity collar that adjusts auction quantities in response to privately held inventories of unused allowances. In fact, in some specifications, the quantity collar performed worse than no stabilization policy at all. The experiment implemented a specific set of structural elements, and extrapolation to other settings should be done with caution. Nevertheless, an examination of the observed behavioral patterns and deviations from optimal behavior suggests that a price collar has an important (although perhaps not exclusive) role to play in constructing an effective market stability reserve policy.

Key Words: EU Emissions Trading System, market stability reserve, price collar, allowance prices, emissions allowances

Contents

Price and Quantity "Collars" for Stabilizing Emissions Allowance Prices: An Experimental Analysis of the EU ETS Market Stability Reserve

Charles A. Holt and William Shobe [*]

I. Introduction

Setting the appropriate path for releasing allowances under a cap-and-trade program such as the European Union Emissions Trading System (EU ETS) is an important regulatory decision with implications for long-term cost-effectiveness. Much depends on the behavior of market participants after the emissions allowances are released into the "wild" of the marketplace. For a stock pollutant like greenhouse gas emissions, the regulator must decide both the long-run, cumulative emissions and the amount to make available in any period. Naturally, if the cumulative supply of allowances is too high, then once this information becomes available to the market, allowance prices will fall too low.[1] In cap-and-trade programs that have been used to date, there appears to be some tendency for emission caps to be set too high, at least early on, so that compliance has been achieved at prices well below those anticipated.

The timing of the release of a given long-run supply of allowances has also long been recognized as a critical design issue for cap-and-trade programs. The concerns about excess price variability due to the short-run fixed supply of allowances are easily addressed for stock pollutants by "frontloading" allowances (selling or granting them some periods before they are needed for compliance) and allowing emitters to bank unused allowances into future compliance periods. Economic theory provides a set of market regularity conditions under which the time path of the frontloaded release should not matter. Essentially, these conditions boil down to the

---

[1] This would occur if the supply of allowances is set higher than the level that would equate the marginal abatement cost with the marginal benefit of reducing emissions.

absence of borrowing constraints, fully informed market participants behaving competitively, and the existence of liquid futures markets. Salant (2015) showed that under these assumptions, if the design of the regulation provides for a minimum effective level of frontloading, then how the remaining allowances are allocated over time will not affect market performance, and emission and abatement decisions will take place along the dynamically optimal path.

If these regularity conditions are violated, then even if the aggregate cap is set correctly, the time path of the release of allowances to the market will matter. It is possible that the trajectory of allowance availability could result in a price that is, in early periods, lower (or higher) than the social optimum, which leads to abatement efforts that are sub-optimally slow (fast), raising the costs of achieving emissions reductions. Where this occurs, policymakers can adjust the regulatory environment to improve market performance so that it better meets the regularity conditions. Alternatively, they can explicitly change the time path of allowance availability or place limits on the range of possible prices through a *price collar* (price ceiling and floor). Whether a given approach, or combination of approaches, improves outcomes depends on the particular circumstances that give rise to the suboptimal abatement path.

The EU ETS is currently faced with a large stock of unused allowances that approximately equals the total yearly emissions covered by the scheme. This situation has resulted in lower-than-anticipated allowance prices, leading to concerns that current levels of emissions reduction effort and investment are too low (European Commission, 2012).[2] The European Commission (EC) has already "backloaded" some allowances by shifting the availability of a large block of allowances from 2014–2016 to 2018–2020, in the hope of supporting higher current prices at the expense of lower prices when these allowances in reserve are released (European Commission, 2014b). The EC is currently considering following this short-term fix with a longer-term policy based on automatic adjustments to preannounced auction quantities that are triggered by abnormally high or low privately held banks of unused allowances. Given the current large private bank of unused allowances, the EU's proposed *quantity collar* would begin by reducing auction quantities, with the unsold allowances being deposited into a market stability reserve (MSR) that could later be used to increase auction quantities in tight conditions with low privately held banks of allowances.

---

[2] See references and discussion in Burtraw et al. (2014).

The EC appears to have based its quantity collar MSR design on two potential deviations of emissions markets from the theoretical benchmark: myopic economic agents and borrowing constraints (European Commission, 2014b, p. 16; 2014a). There is considerable disagreement about how important these market imperfections are relative to other factors that might be driving the current low price in the EU ETS, factors such as overallocation of allowances, regulatory risk, and leakage of emissions from under the EU emissions cap.[3]

This paper uses laboratory experiments to compare the performance of a quantity collar like the one under consideration by the EC with the types of price floors and caps that have been implemented in US regional cap-and-trade programs. We assess how the different policy instruments affect production, net social surplus, and the variation of market price from optimal levels. As measured by these criteria, the MSR sessions do not compare favorably with those of a price collar sessions or even sessions with no policy at all. This surprising finding is at variance with some of the results of computer simulations that have been used to evaluate the performance of the quantity MSR. These differences are discussed in the final section, which also contains some proposed follow-up experiments. The next section describes some of the institutional details of MSR policies to be considered. The third section presents the structure of the emissions markets to be used in the experiments, along with some theoretical predictions associated with extreme cases of myopic behavior and perfect foresight. The laboratory procedures and results are summarized in the fourth and fifth sections, respectively.

II. Price Containment in Emissions Allowance Auctions

Since Rubin (1996) first characterized the intertemporal behavior of emission markets, scholars have paid considerable attention to the optimal path of the cap and of resulting prices. With full banking and borrowing of allowances, the pattern of allowance release by the regulator would not matter if the market meets the standard regularity conditions noted above. But banking without borrowing implies an asymmetry in how market participants treat the present and the

---

[3] Before deciding on the remedy, the regulator needs to consider the cause of the low prices because the market will respond differently depending on the particular reason for the low prices. In the case of an overallocation of allowances, possibly due to slower economic growth than was anticipated when the cap level was set, a quantity MSR would not increase the price unless there was also some borrowing constraint, myopia, or other inefficiency in the market, or unless the policy would be expected to permanently reduce the aggregate intertemporal cap. If the low prices are due to uncertainty over the regulator's commitment to maintaining allowance scarcity, then a different pattern of allowance would be appropriate.

future. For a uniformly mixed, stock pollutant such and greenhouse gases, this asymmetry provides a theoretical justification for a regulatory "frontloading" of allowances—that is, making extra allowances available in advance of the time when they are expected to be needed to cover production. In planning for future allowance needs, firms will generally choose to hold a private bank of allowances to hedge against price risk and, if the cap on emission is declining, in anticipation of higher future prices.

As noted in the introduction, allocating allowances in advance of need facilitates intertemporal arbitrage and cost-effective abatement investment trajectories, but there are conditions under which it is possible to frontload too many or too few allowances. Taschini et al. (2014) demonstrated that, in the EU ETS context, if agents are myopic, then too much frontloading of allowances will result in prices that are too low early on and too high later, leading to higher aggregate abatement costs. Neuhoff et al. (2012) explored the case where compliance entities have lower discount rates for holding allowances than do other participants in the market, possibly due to institutional constraints on their ability to accumulate a private bank.[4]

Fell (2015) used a stochastic, dynamic optimization model to show how the constraint on borrowing allowances from future periods may push prices away from the optimal path. In this case, efforts to lower an overallocated cap, if done without sufficient frontloading, can cause suboptimal price paths to occur when a period of high allowance demand exhausts the private bank. This problem can be reduced by increasing the amount of frontloading or using an allowance reserve.[5] These studies clearly demonstrate that the timing of allocations may matter

---

[4] Based on evidence from interviews with market participants done by Betz et al. (2015) with self-reported discount rates on the order of 3% for firms with compliance obligations and of 15% for other holders of allowances, Neuhoff et al. (2012) showed that, where investors are divided into these two distinct groups, frontloading that is larger than the limited hedging needs of compliance entities may lead to prices reflecting the higher discount rate of speculators and may be substantially lower than the long-run optimal price. Then frontloading too many allowances can result in allowance prices that are sub-optimally low early in the program. Salant (2015) questioned whether the actual experience in the EU ETS market has been consistent with the "hedging corridor" model. He argues that the actual performance of the prices of allowances in the EU ETS has had more to do with market participant perceptions of regulatory risk than with a hedging corridor.

[5] In Fell's models, the price collar reserve mechanism tended to outperform the quantity mechanism of the type under consideration by the EC. In an allowance trading model with uncertainty over future regulatory policy, the pattern of quantity adjustments may be quite different from those now being proposed by the EU; that is, allowances should be retired from later years while maintaining frontloading to avoid having traders hit binding constraints on banking (Salant, 2015).

in the presence of market imperfections. This means that the regulator needs to consider not only the total emissions allowed by the cap, but also whether the time path of release supports the cost effective emission control. The standard presumption that the cap should start fairly slack and then decline steadily over time must be discarded in favor of a more detailed investigation of how market actors actually behave toward a cap policy that will gradually unfold over the next 30 or 50 years.[6]

There are several alternatives to controlling the time path of allocation. Foremost among these would be working to improve the efficiency characteristics of the market itself: lowering regulatory risk, reducing market frictions, and facilitating intertemporal contracting. Progress in these dimensions works to reduce the need for explicitly engineering the allocation path. A second approach, the price collar, is now in use in at least two substantial cap-and-trade programs. A price collar is used in conjunction with a compliance reserve that absorbs unsold allowances in low demand conditions and releases allowances as needed to enforce an upper limit on prices. A price collar can serve to limit longer-term deviations from the optimal price path, which may arise due to market imperfections including regulatory risk, myopia, and limited forward contracting for allowances.

The Regional Greenhouse Gas Initiative (RGGI) program includes a reserve price for the auction that puts a lower limit on permissible bids and serves as a price floor.[7] During the first four years of its operation, RGGI faced a large surplus of allowances due to the concurrent recession and the contemporaneous large drop in natural gas prices. The use of a reserve price resulted in significant quantities of unsold allowances in some early RGGI auctions, which helped maintain auction revenues and prevent an even larger oversupply of allowances.[8] A price floor puts a horizontal segment into the auction supply curve, so that (in theory) quantity adjustments are used when demand is low. The RGGI program now also features a price ceiling

---

[6] The dependence of the optimal pattern of allowance allocations on the characteristics of the market environment greatly complicates not only the design of a given cap-and-trade program but also the prospects for linking disparate programs. If programs are linked, then the time path of allowance availability will be the sum of the allocation paths across all linked programs. No one program could determine the availability of allowances. Thus, a policy of trying to explicitly set the allocation path, while a possible approach for one program, is increasingly less tenable as the number of linked programs grows.

[7] A discussion of the RGGI auction design may be found in Holt et al. (2007).

[8] Some observers have maintained that the whole RGGI program would have collapsed in the absence of a price floor in its early years. Dallas Burtraw from Resources for the Future made this point at the October 2014 MSR Workshop held at DIW in Berlin.

and a regulator's reserve stock of allowances that can be added to the auction quantity to prevent the auction closing price from rising above the ceiling. Just as a price floor adds a horizontal segment to the left of the vertical auction supply curve, the "soft" price ceiling adds a horizontal segment on the right side, with a width determined by the size of the market stability reserve.[9] Together, the floor and ceiling constitute a price collar that limits extreme price variability via in-auction quantity adjustments, while preserving a range in which price signals can guide economic decisions.[10] Such a system could be self-sustaining if unsold allowances at the floor are added to the regulator's reserve, and allowances added at the ceiling are taken out of the reserve stock.[11]

The EC has recently proposed an alternative mechanism for adjusting allowance availability in response to a non-optimal price path (European Commission, 2014a). The quantity collar being considered by the EC implements changes in auction quantities that are triggered by tightness or looseness in allowance markets, as measured by the total "bank" of unused allowances at a point in time. If the bank falls below the lower trigger point, a pre-specified quantity of allowances would be added to a subsequent auction. Conversely, if the bank rises above the upper trigger point, then a subsequent auction quantity would be reduced by a percentage of the total bank of unused allowances. A rationale given for using a quantity collar is that, with myopic agents or agents subject to hedging restrictions, quantity adjustments will work to raise exceptionally low prices and lower exceptionally high ones. This process might result in less variability in the allowance prices than would be experienced with an auction supply that does not change in response to the size of the bank.[12] A quantity collar eliminates the need to

---

[9] Perkis et al. (2014) used experimental methods to compare the performance of "soft" ceilings with limited containment reserves and "hard" ceilings with unlimited containment reserves.

[10] Burtraw et al. (2014) presented a strong argument in favor of a price collar as an alternative to other types of policies for stabilizing emissions allowance prices.

[11] An alternative to an in-auction price cap is to offer extra allowances from the reserve for sale at pre-specified prices after the close of the auction (Shobe et al., 2014). This post-auction sale procedure is used in the Western Climate Initiative (WCI) cap-and-trade program, comprising California and Quebec at this time. The price floor is implemented with an auction reserve price, with unsold allowances added to a base regulator's reserve amount. The price ceiling has a stock of allowances divided into three price tiers. These allowances are available to buyers at prices well above the expected market price, but buffer the market against transient shocks. The post-auction sale is implemented with a rationing rule for allocations of excess demand between different price tiers. The potential adverse effects of this rule are evaluated in an experiment reported by Bodsky et al. (2012).

[12] See Shobe et al. (2010) for a discussion of auctioning allowances under a loose cap.

specify an acceptable price range for allowances.[13] Nevertheless, one crucial aspect of the proposed EU program is the need to pre-specify the upper and lower trigger points on the stock of unused allowances. If the limits are sufficiently loose, then they might never bind, and if they are tight, they might produce unwanted price cycles (Trotignon et al., 2014).

An important difference in implementation is that the price collar induces contemporaneous quantity adjustments, whereas the quantity collar produces an adjustment in subsequent auctions. This change would occur with a lag, since time is needed to obtain accurate measures of the bank of unused allowances. Price and quantity collars differ in terms of timing and effect, so it would be possible to implement a hybrid of both methods by using a common reserve stock.

Any price containment policy, regardless of how it is structured, is likely to affect the use-or-bank strategies of market participants in a manner that is difficult to model. Such strategies are made in the absence of perfect foresight about demand, cost, and even regulatory conditions. Given the limitations of modeling and simulating dynamic behavior and the associated trader forecasts on which it is based, this study makes use of laboratory experiments.[14] Such experiments have been extensively used to test and evaluate auction designs, such as for RGGI and California emissions allowances, FCC broadcast spectrum, irrigation permits, and "toxic" mortgage backed securities.[15] The next section describes the basic experiment design and the associated static and dynamic price predictions in the absence of price containment mechanisms.

---

[13] Setting such a range appears to face strong political barriers in the current EU context, since a narrow price collar could be construed as a tax with a more stringent approval process.

[14] Cason (1995) uses experiments to evaluate the effects of an early emissions auction that was run by the US Environmental Protection Agency. See Cason (2010) and Friesen and Gangadharan (2013) for recent surveys of experiments motivated by issues in emissions allowance trading.

[15] The design for the RGGI uniform price auctions was based on extensive experimentation reported in Holt et al. (2007). The uniform price auction proved to be more resilient to explicit and tacit collusion than two versions of an ascending "clock" auction (Burtraw et al., 2009). A clock auction was used to sell nitrous oxide allowances for the state of Virginia, after being tested in the lab (Porter et al., 2009). The FCC's first major auction that incorporated bids on packages of licenses was designed and tested with laboratory experiments (Goeree and Holt, 2010). The US Treasury's proposed TARP auction (for buying mortgage backed securities from banks in distress) was designed by experimental economists J. Goeree, C. Holt, and C. Plott, working with O. Armantier at the New York Federal Reserve bank. This auction was canceled when the Treasury decided to proceed quickly with direct asset purchases from banks during the window of opportunity prior to the 2008 election. The basic TARP auction design, however, was tested subsequently by Armantier et al. (2013). See Cummings et al. (2004) for a report on an auction for irrigation reductions conducted during a major Georgia drought, after being tested with laboratory experiments.

III. An Experiment Design with Binary Demand Shocks: Static and Dynamic Predictions

These experiments test how different emission market price collar mechanisms affect market outcomes. In order to focus on time-shifting of allowance demand, we characterize the abatement decision as either the switching of production from high emitters to low emitters or cutting back on output. In particular, subjects with producer roles are designated as being either "low users," who require only a single allowance for each product unit produced, or "high users," who emit twice as much and require two allowances per unit. All subjects have a number of "capacity units," each of which can be used to produce a unit of a product to be sold at an exogenous random price. Capacity units have costs of operation that are randomly determined. Even though low users in the experiment tend to have higher costs on average, the relevant marginal costs are determined by taking the sum of the production cost and the price of needed allowances, which is higher for high users.

Each experimental session is a sequence of "periods," where each period starts with announcements to subjects about state variables that depend on what happened in previous periods and then proceeds with an auction of a specified number of allowances (referred to as "permits" in the experiment), using a uniform price (highest rejected bid) procedure that is standard in the EU and in regional US allowance markets.[16] For a subset of the treatments, the auction is followed by a spot market in which participants can buy and/or sell permits by submitting bid and ask limit orders to a market maker. Participants then decide which capacity units to operate for production of a product that is then sold at a fixed price.

Participants must cover each unit of output with the required number of permits (1 or 2). Those in deficit must pay a substantial fine ($20) plus any deficits are carried forward to future periods until covered by a permit. Unused permits are banked for use in future periods. Auction quantities are reduced over time in a preannounced manner, but preannounced auction quantities can be adjusted by the regulator using a price collar or quantity collar instrument. Subjects earn money based on the difference between prices of product units sold and production costs,

---

[16] See Lopomo et al. (2011) for a thoughtful evaluation of alternative auction procedures for allowance auctions, and of uniform price auctions in particular. In a multiunit uniform price auction, bids need not match permit values because bidders may have an incentive to bid strategically in an effort to alter the clearing price. There are mechanisms, such as the Vickrey multiunit auction, that do not have this property, but such auctions are not commonly used. See Ausubel and Milgrom (2005) for a thorough discussion of the theoretical properties of the Vickrey auction and of practical considerations (complexity, revenue generation) that limit its use in practice.

including the purchase of required permits. In sessions with spot markets, subjects can also earn money by trading in spot markets.

The product market is subject to significant demand shocks, randomly determined "high demand" and "low demand" periods with high and low output prices respectively. Each period, the participants see the output price prior to the current auction.

The current value of a permit used in production is determined by taking the difference between the product price and the production cost for the capacity unit being used, and then dividing that difference by the number of permits required (2 for high users, 1 for low users). In a particular period, the permit values for all participants can be ranked from high to low in order to generate a demand function for permits, which can be crossed with a supply function that is vertical at the auction quantity. This static price prediction would fluctuate up and down due to random demand and cost shocks. Given our multi-period sessions, permit values depend on long-run considerations, with the socially optimal abatement path equating discounted marginal abatement costs over time as well as across firms. In a model with no discounting, an ex post prediction of the dynamic price can be obtained by using the sequence of realized costs and product prices to calculate permit values for each capacity unit in each period, and to rank these values for all periods from high to low and cross them with the total supply for all auctions combined. Permits with values above this cutoff should (with perfect foresight) be used in production, and those with values below the cutoff should be banked. These static and dynamic efficiency measures can then be compared with the auction and spot market prices observed in the experiment.

A straightforward derivation of ex ante efficiency predictions can be based on the maximization of a surplus measure, subject to a constraint on the total number of permits. This maximization requires that the "use values" of permits be equalized at the margin—that is, equality of the differences between the price of the product being produced and the marginal cost of production, with each such difference being divided by the number of permits required for

each unit of production (2 for high users and 1 for low users).[17] In other words, the marginal valuation (price minus marginal cost divided by the required number of permits) equals a constant λ for all periods (the Lagrange multiplier), which turns out to be the dynamic permit price prediction that could be obtained from a standard analysis of supply and demand.[18]

The production costs used in the experiment are drawn from uniform distributions on specified intervals, which are easy to explain to subjects and which result in linear marginal costs. Let K_L denote the total number of capacity units for all low users combined (e.g., the number of low users times the number of capacity units each). If costs for low emitters are independently drawn from a uniform distribution on the interval $[a_L, b_L]$, the *marginal* cost of producing Q_L units can be approximated by the function $C'(Q_L) = a_L + Q_L (b_L - a_L)/K_L$, with a marginal cost of a_L for a quantity of 0 and a marginal cost of b_L at full capacity with $Q_L = K_L$. If the exogenous product price (or its expected value) is represented by P, then the marginal profit for producing Q_L units is the difference between product price and marginal cost: $P - C'(Q_L)$. Each unit of output for a low user requires 1 permit, so $x_L = Q_L$ and the inverse demand for permits is $p_x = P - C'(x_L) = P - a_L - x_L(b_L - a_L)/K_L$, where p_x is the price of permits and x_L is the quantity of permits demanded by low users. This function can be solved for x_L to obtain the low users' demand for permits as a function of the permit price p_x. The derivation of the permit demand for high users is analogous, with appropriate changes in notation.[19]

Total production capacity was 20 units for each producer type: $K_L = K_H = 20$. The cost range was specified to be [0, 30] for high users and [10, 30] for low users, so $a_H = 0$, $b_H = 30$, a_L

---

[17] Let the value of the product be a function $S(Q_{Lt} + Q_{Ht})$, where Q_{Lt} and Q_{Ht} denote the aggregate production of low and high users, respectively. This analysis is based on a partial equilibrium assumption that the product price P_t in period t represents the marginal social value of the product: $S'(Q_{Lt} + Q_{Ht}) = P_t$. Let x_{Lt} and x_{Ht} denote total permit use by low and high users. Since each unit produced requires one permit for low users and two for high users, $Q_{Lt} = x_{Lt}$ and $Q_{Ht} = x_{Ht}/2$. Thus the surplus difference in period t can be written as $S(x_{Lt} + x_{Ht}/2) - C_L(x_{Lt}) - C_H(x_{Ht}/2)$, where C_L and C_H are the total costs associated with production by low and high users. The Lagrangian is the sum over t of these surplus value differences, with a constraint that the sum of all permit use quantities for all periods equals the sum of auction quantities. The partial derivative of this Lagrangian with respect to x_{Lt} yields $P_t - C'_L(x_{Lt}) = \lambda$, where λ is the Lagrange multiplier for the constraint on the total number of permits. The analogous condition for high users is $[P_t - C'_H(x_{Ht}/2)]/2 = \lambda$, where the divisions by 2 are caused by the fact that the production of high users is half of the number of permits used.

[18] With discounting, the present values of price/marginal cost differences would be equated to λ, so that undiscounted price/marginal cost differences would be equated to $(1+r)^t \lambda$, which is the standard result that the price of permits (cutoff value of permits) rises at the rate of interest.

[19] The permit demand functions for low and high users are $x_L = (P - a_L - p_x)K_L /(b_L - a_L)$ and $x_H = (P - a_H - 2p_x)2K_H /(b_H - a_H)$.

= 10, and b_L = 30. If the exogenous auction quantity is 35, these parameters can be used with the permit demand functions for each user type and summed to obtain the demand for permits in a single period as a function of the product price P:

$$x_L + x_H = P(7/3) - 10 - (11/3)p_x = 215/3 - (11/3)p_x \text{ if } P = 35. \tag{1}$$

If the product price is 35, then the demand on the far right side of (1) can be equated to the auction quantity of 35 to determine the equilibrium permit price, which is 10. Then the demand functions for each user type can be used to show that low users purchase 15 permits, operating 15 of their 20 capacity units, and high users purchase 20 permits in total, operating only 10 of their 20 capacity units. Thus the equilibrium permit price of 10 results in lower capacity utilization for high users. Finally, each producer is required to produce with a minimum of one of their four capacity units, to cover "must-serve" sales.

Next, consider what changes when the product market price switches randomly between high and low levels of 40 and 30, with an average of 35. To graph the demand for permits with the permit price on the vertical axis, the left-side equality in (1) is solved for p_x as a function of the product price, P, and the total demand for permits for both user types, x, to obtain the inverse demand for permits: $p_x = (3/11)[(7/3)P - 10 - x]$. This demand function is graphed in Figure 1 for a high product price of 40 (upper downward-sloping gray line), a low price of 30 (lower downward-sloping gray line), and an average price of 35 (dotted line in between). At the average product price, the demand for permits crosses the vertical supply segment at a permit price of 10, which is the dynamic price prediction (horizontal dashed line).

Of course, it is not optimal to use exactly 35 permits in each period. On average, optimal timing of production involves using about 47 permits in high demand periods and 23 in low demand periods, as indicated by the diamond-shaped marks, with an average use of 35 permits. The diamond on the left side of the figure has low users operating at half capacity and high users are at one-third capacity, with both users banking permits for future use in high demand periods. At the right-hand diamond, low users are operating at full capacity and high users are at two-thirds capacity. The basic structure stays the same from period to period, so the analysis of the Walrasian equilibrium for a single period extends to the entire sequence.

**Figure 1. Permit Price Predictions with Demand Shifts
(Product Prices of $30, $35, and $40)**
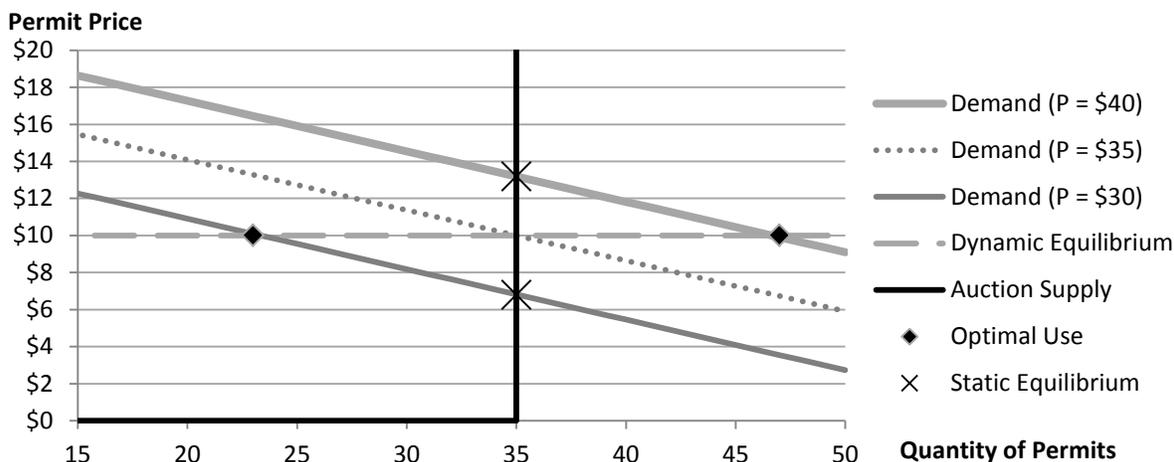
These dynamic predictions contrast sharply with the static equilibrium (with no banking), determined by the intersection of the relevant permit demand function with the vertical auction supply. In a high demand period with a product price of $40, this intersection is at a quantity of 35 and a permit price of about $13, as shown by the "x" mark in the upper part of the vertical supply function in Figure 1. In a low demand period with a product price of $30, this intersection is at a price about $7, again at a quantity of 35, as indicated by the lower "x" mark in the figure. The static equilibrium permit price adjusts sharply to changes in current demand, while the dynamic equilibrium price is constant, and only permits with use values at $10 or above are used, which causes adjustment in the quantity dimension.

IV. Experiment Procedures

Several sessions were run as software tests, and some insights for design improvements emerged. Most notably, the number of auction periods was increased from 12 in the pilot sessions to 18 (series A) or 30 (series B) in order to evaluate price containment policies in both loose and tight environments and in the transition phases. The parameters used in these two series are shown in Table 1.

Series A had 3 treatments (3 sessions each): no collar, price collar, and quantity collar. A The same treatments were used in Series B, along with an additional treatment that used a high reserve price collar. Each session had 10 participants (5 low users and 5 high users), for a total of

210 participants, all students at the University of Virginia. All sessions had the product price being equally likely to be $30 or $40, and with a theoretical dynamic permit price of $10.

Initial cash endowments were $50 for low users and $100 for high users in both series. Each person also received an initial endowment of permits; the total endowment was raised from 45 in series A to 135 in series B,[20] to generate a larger initial surplus supply, that is, a greater level of frontloading. Auction quantities started at 41 (series A) or 45 (series B) and declined by 1 in each successive auction. These parameters were selected to ensure that the total number of permits allocated in all periods was equal to 35 (as in Figure 1) times the total number of periods, which would keep the dynamic price prediction constant. In fact, the dynamic price prediction for the realized random cost and product price draws turned out to be $10.25 for both series A and series B, as shown in Table 1.[21]

In series A, each auction was followed by a spot market in which participants could buy and/or sell permits from others at a common "call price" determined by the intersection of bid and ask arrays.[22] The increase to 30 periods for series B sessions was made possible by eliminating spot markets, which were not very active in the series A sessions. Subjects in series B were not permitted to submit bids in a given auction that exceeded their current cash balance plus the sum of current-period use values for their capacity units. A final difference was that the subjects were told the probability of a high demand period (0.5) in series B, whereas the instructions did not reveal this probability in series A. This last change was intended both to reduce cognitive demands and to lessen extraneous learning effects that were observed in some early sessions.

---

[20] These initial endowments are included in calculations of the aggregate session cap and do not reflect a change in the session level scarcity of permits.

[21] This dynamic price calculation for each series was done by ranking the permit use values, as determined by the product price and capacity unit cost realizations, and then crossing this array with the sum of all auction quantities for all periods, plus the initial endowment.

[22] See Davis and Holt (1993, chapters 3 and 4) for a survey of experimental research on the efficiency properties of call markets.

Table 1. Experiment Design

	Series A	Series B
Number of sessions	9	12
Subjects per session	10 (5 low users, 5 high users)	10 (5 low users, 5 high users)
Number of periods per session	18	30
Initial cap	41	45
Initial permit endowment	3 low, 6 high	9 low, 18 high
Initial cash endowment	50	100
Total permits (incl. endowment)	630	1050
Unconstrained demand	900	1500
Initial regulator's reserve	20	30
Bank trigger quantities (low/high)	30/55	65/90
MSR injection	7 (~12% of 55)	11 (~12% of 90)
MSR retirement rate	12% of private bank	12% of private bank
Dynamic price	10.25	10.25
Spot market	yes	No
Price collar	8 to 12	8 to 12 or 9.5 to 13.50
Treatments	No collar (3 sessions)	No collar (3 sessions)
	Price collar 8-12 (3 sessions)	Price collar 8-12 (3 sessions)
	Quantity collar (3 sessions)	Quantity collar (3 sessions)
		High price collar (3 sessions)

The treatments used for each series are in the bottom panel of Table 1. Three treatments were common to both series: a baseline control without any price containment policy, a price collar, and a quantity collar. The price collar was selected to be symmetric around $10, with a floor of $8 and a "soft" price cap of $12. Any unsold allowances at the price floor are deposited into the regulator's price containment reserve, which begins at a level of 20 permits. Conversely, permits are withdrawn from the reserve to prevent the auction clearing price from rising above the price cap, to the extent possible given the current reserve. For the quantity collar in series A, the lower and upper trigger points were set at 30 and 55 permits, respectively. The auction quantity would be increased by 7 in the second auction that followed a situation in which the bank of unused permits fell below 30. Conversely, a reduction of 12 percent of the bank of unused permits would occur in the second auction following a breach of the upper quantity point of 55.

Series B had one additional treatment, a high price collar ($9.50–$13.50). The use of a high price collar was motivated in part by the recognition that a regulator would be unlikely to have the information needed to bracket the dynamic price in a symmetric manner. Moreover, the high price collar could increase the possibility of quantity reductions caused by unsold units at a higher price floor. The quantity collar trigger points for series B were increased to 65 (low) and 90 (high) to be more in line with the higher initial free allocation for this series, with an increase of 11 triggered by a breach of the lower limit and a decrease of 12 percent of the privately held bank triggered by a breach of the upper limit.[23]

The regulator's initial price containment reserve for both price and quantity collar treatments was raised from 20 in series A to 30 in series B. All sessions in each series were run with the same sequences of demand and cost realizations in order to maintain balance across treatments, but the random sequences used in series B differed from that used in series A.

Our experimental design does not mimic the hedging collar framework in the computational model used in Neuhoff et al. (2012). We do not impose the types of institutional constraints that place a cap on the hedging demand of one set of subjects but not another. While the participants in our experimental sessions clearly have varying preferences toward the size of their bank, there is not any reason to suspect a clear break in hedging demand as between unconstrained investors and constrained investors. The initial free allocation of allowances to subjects is never high enough to satiate all hedging demand, although, because of the wide variation in hedging activity of subjects, many of them stop accumulating banks well before the aggregate bank reaches the trigger quantity for the MSR mechanism to reduce future auction quantities.

Participants in the experimental sessions were all students at the University of Virginia. Students are recruited through an email invitation to those who have signed up to be part of the general pool of subjects for social science experiments at U.Va. The experiment is couched in generic terms such as "output" and "permits" to avoid context that may affect their responses. Sessions took about 1.5 hours and students earned just over $25 on average. The students are

---

[23] The size of the auction increase amounts, 7 and 11, were chosen to mimic the relative magnitudes of increases and reductions found in the MSR proposal at the time of this writing. They are approximately 12% of the upper trigger amount.

highly motivated, and it is clear from the observed outcomes that they are able to take advantage of available arbitrage opportunities.[24]

V. Experiment Results

We begin with a display of auction prices and other data for representative sessions in series A in order to provide some context for the discussion of overall averages by treatment that follows. Figure 2 shows the sequence of auction clearing prices for two sessions, one without any price containment policy in the top panel, and one with the price collar between $8 and $12 shown in the bottom panel. In each case, the sequence of static price predictions with no banking is plotted as a dashed gray line with a jigsaw pattern that responds to random changes in the product price. The actual auction prices are shown as larger dark dots. In both sessions, the observed auction prices start near the dynamically optimal price of about $10 (horizontal dashed gray line), but prices in the session with no collar fall subsequently into a lower range ($5–$6).

The price collar limits ($8 and $12) were selected to be binding in the sense of preventing sharp movements to static, no-banking predictions determined by the "x" marks on the vertical supply line in Figure 1. The observed price sequence in the bottom panel of Figure 2 did not run up against these bounds (with a couple of exceptions), despite the fact that the static Walrasian predictions were either below the price floor or above the ceiling price in almost all periods. It appears to be the case that forward-looking dynamic behavior with banking softened the predicted price gyrations to a great extent in this session, even without the help of a price collar. There were 9 unsold permits, however, when the final auction in the price collar session closed at the reserve price. Several subjects acquired large banks in early auctions, which raised prices above the reserve price. These people leaked permits into the spot market and used them wisely later in high demand periods. As a result, the auction-clearing prices in the bottom panel of Figure 2 tend to be fairly flat and only slightly above this prediction during the sequence of high demand periods in which auction quantities fall toward the end of the session.

---

[24] There is some evidence that using subject matter experts in a simplified laboratory context can lead to less reliable outcomes. Experts may use rules of thumb from their own work context, which may be inappropriate in the lab. Also, those with knowledge of the subject of the experiment may have opinions about policy outcomes that prejudice their choices in the experiment.

Figure 2. Auction Prices (Dark Dots) and Static Predictions (Gray Dashed Line) with No Collar (Top Panel) and with a Price Collar (Bottom Panel)
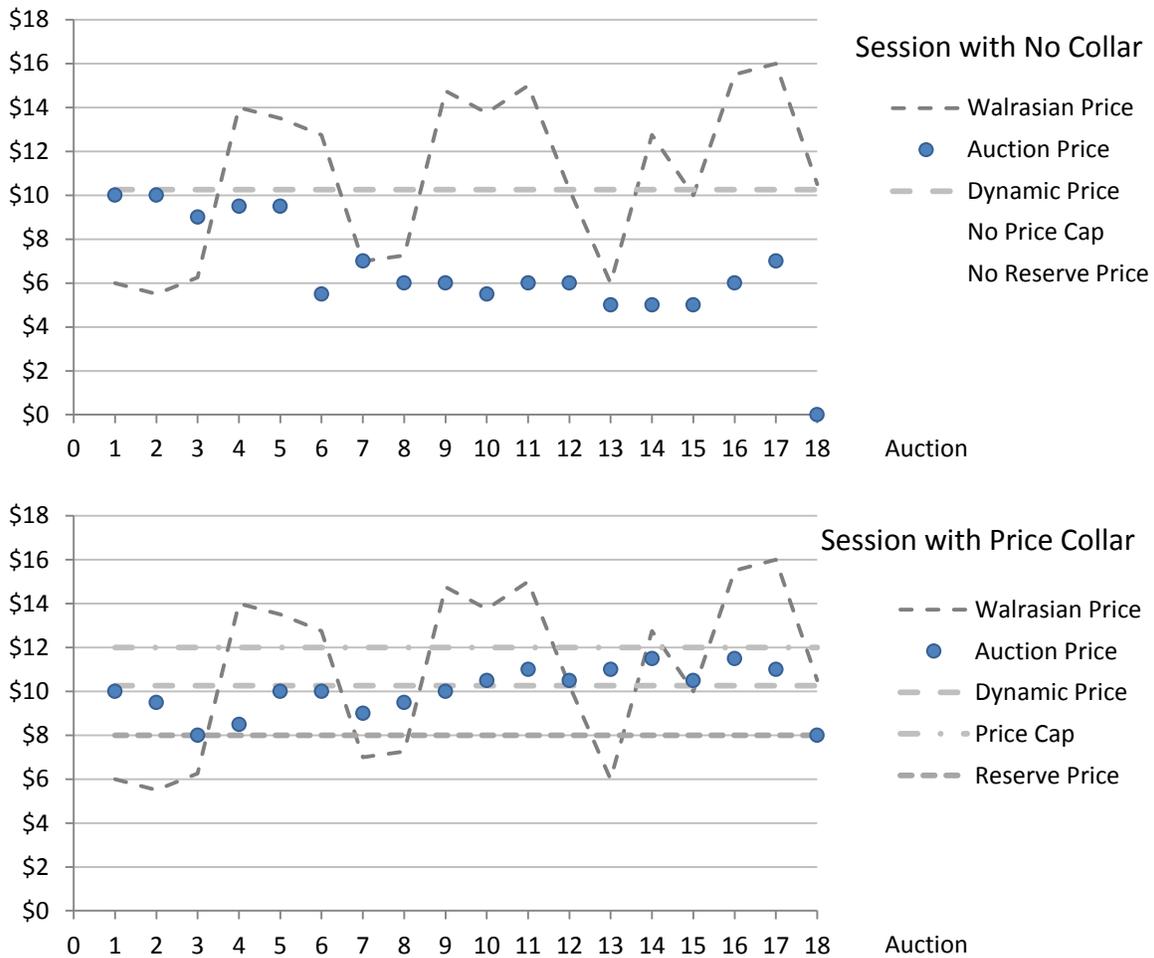
Next, consider Figure 3, in which the top panel shows the comparable auction price series for a session with a quantity collar, with the same sequence of demand shocks and cost draws that were used for the previously discussed sessions. Even without a price floor, the clearing prices in the first three auctions shown in the top panel of Figure 3 are well above no-banking Walrasian predictions that result from high initial auction quantities. There was clearly some strategic, forward-looking behavior, with several people acquiring large banks in early periods. In fact, the total bank of unused permits rose from the initial endowment level of 45 to levels of 59, 79, and 83 prior to the second, third, and fourth auctions. These inventories exceeded the upper trigger point of 55, resulting in auction quantity reductions in auctions 4–6, which were high-price periods with resulting high current-use permit values. The combined effects of these auction quantity reductions and high current product prices probably contributed to the high auction clearing prices (about $14) observed in periods 4–6.

Figure 3. Quantity Collar Session
Top Panel: Auction Prices (Dark Dots) and Static Predictions (Gray Dashed Line)
**Bottom Panel: Banked Permits, Base and Adjusted Auction Quantities**
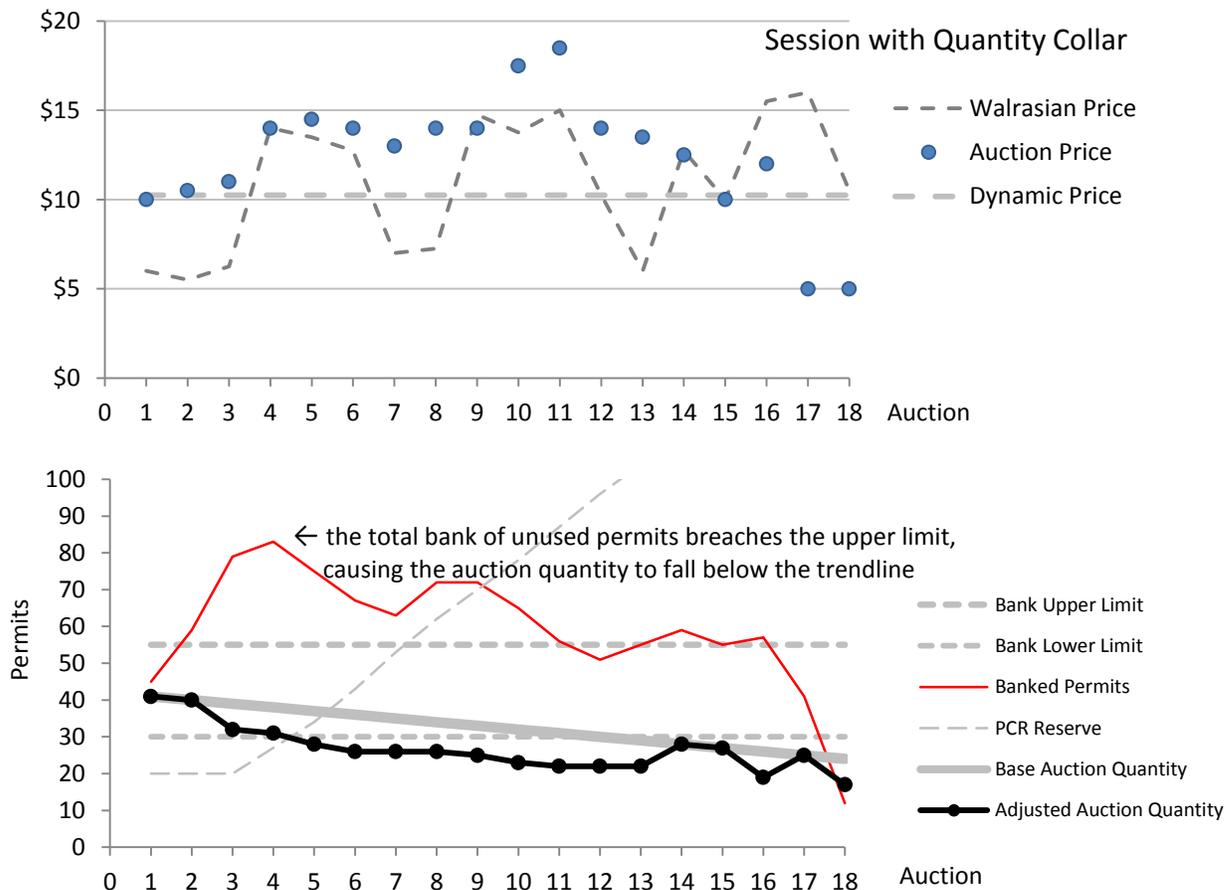
The sequence of quantity collar adjustments is shown in the bottom panel of Figure 3. The downward-sloping thick gray line showing declining *base* auction quantities is overlaid with a line connecting dark dots that tracks *adjusted* auction quantities. As noted above, slack conditions and permit banking in early periods caused the stock of unused permits (thin solid line) to breach the upper limit of 55 after the first period, which caused auction quantities to be reduced in auction 3 and in most auctions thereafter. Although the total bank of unused permits (shown by the thin solid line in the bottom panel) did tend to fall subsequently, it generally stayed above the upper trigger point of 55, resulting in tight conditions for most of the remainder of the session. The regulator's price containment reserve (thin gray dashed line) rose steadily and ended up being more than triple the average auction quantity.

Average Auction Prices

Figure 4 provides a summary comparison of the auction price series, averaged across all sessions within each treatment for each series. The static ("no-banking") price predictions are shown by the thin dark dashed line. These static Walrasian predictions show sharp increases in high demand periods, for example the increase from $6 to $14 in period 4 in the top panel. The dynamic price of $10.25, as determined by the cost and demand realizations, is shown by the horizontal dashed gray line that is the same for all treatments.

Figure 4. Auction Prices Averaged over All Sessions in Series A (Top) and Series B (Bottom) by Treatment
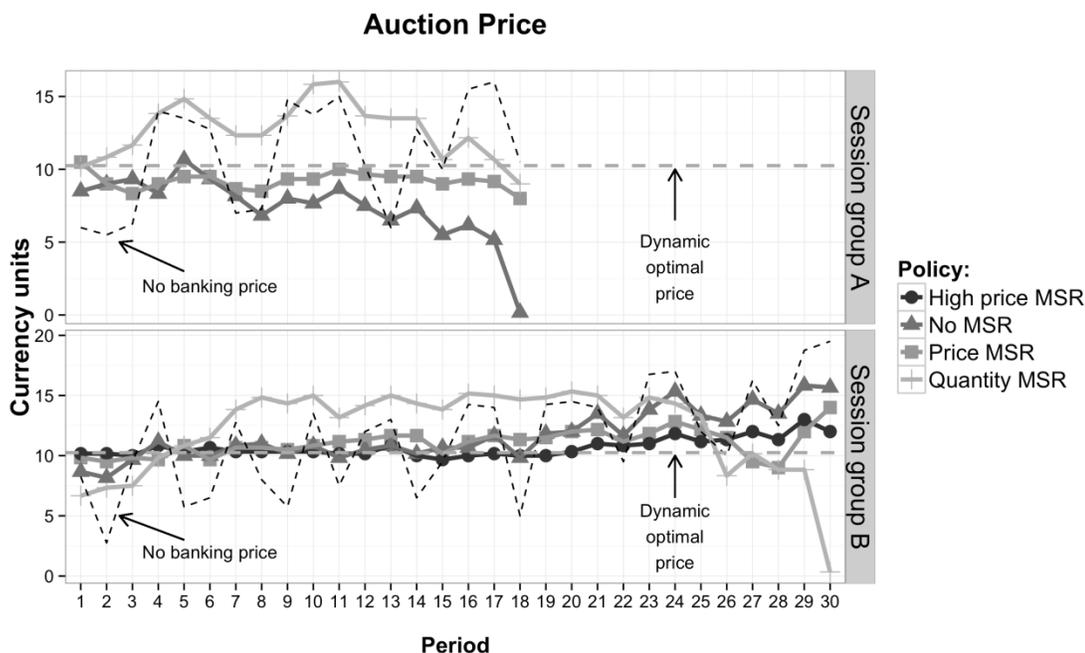
First consider the auction prices in the absence of any price or quantity-based MSR policy (sequence of connected triangles in each panel). These no-collar prices tend to be lower than the other series in the top panel, with average prices generally falling below the $8 level in the second half of the experiment. These low prices are indicative of the effects of the induced "looseness" in the permit markets as a result of the high initial endowments. The opposite case is apparent for the no-collar price in the 30-period sessions in the bottom panel (again shown by connected triangles). The upward spikes in prices in the later periods of the longer sessions with no collar were triggered by the failure of subjects to build a sufficient bank, and this apparent myopic behavior resulted in high prices in the final periods that could not be restrained in the

absence of a price cap in this treatment. With a symmetric price collar, by contrast, prices stay in the $8–$12 band in each series, and do not tend to run up against the boundaries until the final periods. In both panels, the auction price series for the quantity collar (connected plus signs) is higher and somewhat more variable than is observed for the price collar.

Banking

The motivation for the upper limit on the quantity collar is that when the bank of unused permits breaches that limit, as happened in all of the quantity collar sessions, then the reduction in auction quantity would create a tighter situation in which the high banks would fall. Figure 5 shows the sequence of total privately held banks of unused permits, averaged over treatments for each series. The bank sequence is notably lower for the quantity collar in the longer, 30-period sessions, in which auction quantity reductions persisted for many rounds. At this point, we can only speculate about why the auction tightening of the quantity collar did not effectively reduce the bank of privately held permits in series A. Some of those with high banks offered to sell a few permits at high prices in the spot market, which suggests that the observed auction price increases fueled speculative banking. Even those with low banks tended to be cautious and reduce production in order to reduce the risk of penalties or expensive last-minute spot market purchases to cover the compliance obligation of their "must-serve" units.

Firms participating in emission markets have a variety of reasons to hold a bank. Their demand for banked allowances is presumably like that for other inputs and is downward sloping in price. As policy raises the price of banking, the marginal value of banked permits rises. In the absence of perfect arbitrage or if there is regulatory risk, this may raise allowance prices, as desired. But it also raises the cost of using a banked allowance and will change the relative value of other things the firm does. This may lead to unintended consequences.

When the quantity collar drives up the marginal value of banked permits, then the use value of the marginal allowance used in production will be too high. Figure 6 shows actual production, averaged across sessions, for the treatments in both series. During periods where the quantity collar reduces auction quantities, production is much less responsive to changes in the output price, due to the higher opportunity cost of using allowances, as is apparent from the low and relatively flat production series for the quantity MSR series in Figure 6. This pattern of underproduction in high output price periods is especially costly in terms of lost surplus, since there are periods of high production value, as shown by the dashed gray optimal production line. In series B, where there are sustained releases from the reserve in later periods, the pattern of underproduction is less apparent.

Figure 5. Total Banks of Unused Permits Averaged over All Sessions by Series and Treatment: Quantity Collar, Price Collar, and No Collar
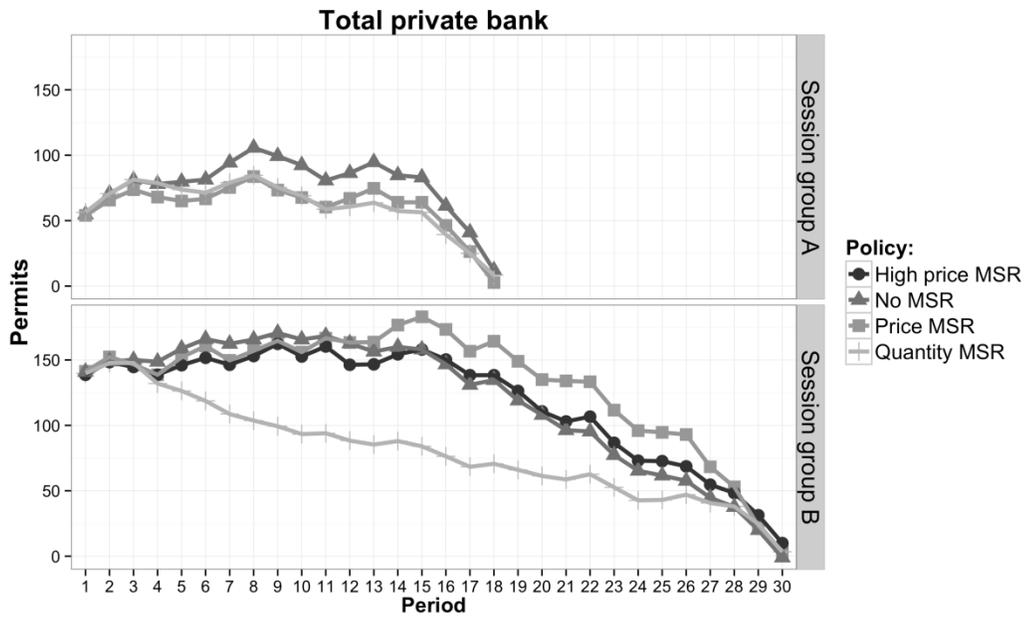
Figure 6. Production Levels by Treatment
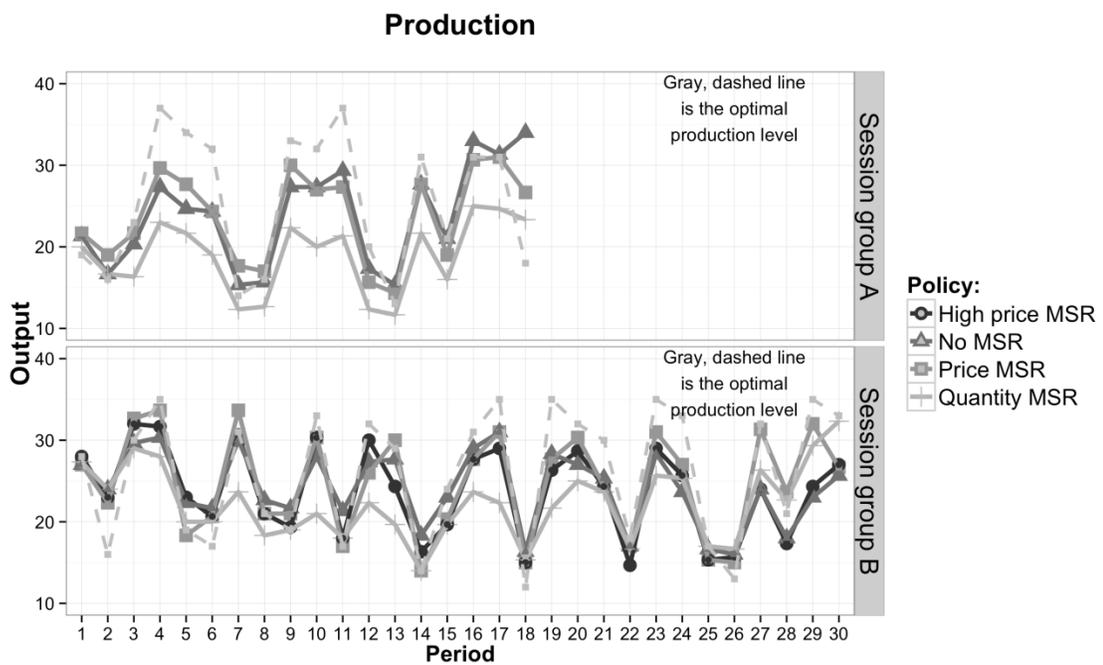
Aggregate Performance Measures

Tables 2 and 3 provide summary information, for the two series respectively, on permits used (left column) and price averages and variances, by treatment (second column). For the 18-period sessions reported in Table 2, the high and somewhat variable auction prices for the quantity collar treatment indicate that this policy is effective at raising the price of emissions permits as intended, but as noted above, the high prices do not, in these sessions, reduce the privately held banks of unused permits. The auction price series for the quantity collar treatments show some tendency to track the static no-banking predictions in Figure 4, so it is not surprising deviations of production from optimal levels (middle column of Tables 2 and 3) tend to be higher for the quantity collar sessions. As a consequence, the efficiency measures for this treatment tend to be lower. Our efficiency measure is the net social surplus from production minus the social cost of any emissions that occur as a result of production ($10.25 for low emitters and $20.50 for high emitters). It is calculated as the sum across all units of production used of output value minus production cost minus the social cost of emissions ($10.25) as a percentage of the maximum that would result from using the dynamic price as a cutoff for deciding whether to use or bank permits. This method for calculating adjusted efficiency implicitly credits the social surplus for any emissions that do not occur because a unit of output is not produced. Efficiency numbers are shown in the far right column of each table.

Table 2. Summary Performance Measures for 18-Period Sessions in Series A

	Total Permits Used	Avg. Auction Price (variance from optimal)	Actual vs. Optimal Production (mean abs % diff.)	Efficiency
Dynamic Optimum:	**630**	**$10.25**	**0**	**100%**
No Collar 1	623	$6.55 (19.28)	15.7	79.3%
No Collar 2	599	$7.83 (16.56)	18.6	80.2%
No Collar 3	623	$7.75 (11.34)	19.5	83.1%
Price Collar 1	621	$10.00 (1.17)	13.4	87.1%
Price Collar 2	625	$8.61 (3.56)	21.8	76.2%
Price Collar 3	612	$9.03 (1.81)	14.0	90.1%
Quantity Collar 1	522	$12.38 (16.17)	24.9	77.5%
Quantity Collar 2	506	$12.53 (8.34)	21.8	79.0%
Quantity Collar 3	472	$13.11 (11.78)	29.5	77.8%

Table 3. Summary Performance Measures for 30-Period Sessions in Series B

	Total Permits Used	Avg. Auction Price (variance from optimal)	Actual vs. Optimal Production (mean abs % diff.)	Efficiency
Dynamic Optimum:	**1050**	**$10.25**	**0**	**100%**
No Collar 1	1058	$12.13 (18.18)	21.8	75.11%
No Collar 2	1049	$11.25 (2.98)	13.5	80.03%
No Collar 3	1046	$11.48 (5.05)	17.3	76.00%
Price Collar 1 $8-$12	1082	$11.32 (3.76)	12.0	86.5%
Price Collar 2 $8-$12	1079	$11.55 (4.33)	16.5	76.5%
Price Collar 3 $8-$12	1068	$10.32 (1.26)	10.9	90.8%
Price Collar 4 $9.50-$13.50	1037	$10.02 (0.81)	14.8	78.1%
Price Collar 5 $9.50-$13.50	1034	$11.98 (4.25)	15.6	83.1%
Price Collar 6 $9.50-$13.50	1000	$10.03 (0.63)	20.4	78.5%
Quantity Collar 1	933	$12.50 (17.55)	18.9	80.7%
Quantity Collar 2	1073	$11.28 (11.38)	20.6	78.3%
Quantity Collar 3	852	$12.65 (25.01)	25.7	71.1%

We use a permutation test, stratified by session series, to test for treatment effects on three key measures of market performance: adjusted efficiency, the difference between production and optimal production, and the variability in price as measured by the mean-squared deviation of the auction price from the dynamically optimal price. In each case, we test the null hypothesis of no effect against the (two-tailed) alternative that there is a difference in outcomes across treatments. The outcomes of these tests are summarized in results 1–3, which follow.

The most striking aspect of the efficiency numbers in Table 2 is that all three sessions with no MSR policy have higher efficiencies than the three sessions with the quantity collar. The price collar treatment also has higher efficiencies than the quantity collar treatment, although there is some overlap in terms of adjusted efficiencies in the far right column. These general patterns also emerge from a consideration of the efficiency measures for series B, shown in Table 3, although there is some overlap in the comparisons.

Result 1 (Efficiency): Efficiency is significantly higher with the price collar than with the quantity collar. There is no significant difference in efficiency between the quantity collar and no collar.

Support: As shown in the top row of Table 4, the average adjusted efficiency is over 5 percentage points higher with a price collar than with a quantity collar. The null hypothesis of no effect is rejected with a p value of about 0.06. (The p-values reported in the table are for a 2-tailed test, with permutations of session efficiency measures within each series, with * indicating significance at the 0.10 level.) There is no significant difference between the quantity collar and no policy at all. Finally, while there appears to be some efficiency gain (about 4 percentage points) from using the price collar relative to no policy, this fails to reach the 10 percent level of significance in a two-tailed test. Further tests on the efficacy of the price collar relative to no policy are warranted. Subjects in the experiments are able to accomplish substantial smoothing even in the no-policy case. In the case of less forward-looking subjects, the performance of the price collar would likely improve relative to the alternative of no policy.

Table 4. Stratified Permutation Test of Adjusted Efficiency (E) Differences

Comparison	E Price – E Qty.	E Price – E No Policy	E No Policy – E Qty.
Difference (p-value)	5.46 (0.057*)	4.01 (0.112)	1.45 (0.400)

Significance level: * - 10%, ** - 5%, *** - 1%.

As shown in Tables 2 and 3, the quantity MSR effectively raises the price of emissions allowances above the price that occurs either with no policy or even with price collars. Figure 4 shows the mechanism through which this effect occurs. Under the quantity MSR, producers are less responsive to realizations of high output prices; in other words, they do not increase output as much during high output price periods in the quantity MSR treatment as they do in other treatments. This leads to efficiency losses during high output price periods relative to the optimum as producers hold on to their banked allowances rather than increase production during these periods.

Result 2 (Optimal Production): The price collar and the no-collar treatments yield significantly lower deviations of observed from optimal production than is the case for a quantity collar.

Support: Once again, as shown in Table 5, we easily reject the null hypothesis of no difference between the price collar and the quantity collar, in this case, in the mean absolute deviation of actual production from optimal production across treatments. There is also a significant difference, although somewhat smaller in magnitude, between the quantity collar and no policy.

With the quantity collar, producers are not as responsive to higher output prices as are producers in the other two policy treatments. Producers appear to place more weight at the margin on banked allowances, which induces them to forgo some profitable production opportunities, especially during periods of high demand for output.

Table 5. Stratified Permutation Test of Deviations between
Actual and Optimal Production (PD)

Comparison	PD Price – PD Qty.	PD Price – PD No Policy	PD No Policy – PD Qty.
Difference(p-value)	-8.24 (0.002***)	-2.31 (0.283)	-5.94 (0.015**)

Significance level: * - 10%, ** - 5%, *** - 1%.

Our third hypothesis addresses the issue of price variability. Emissions market theory suggests that, in well-behaved market settings, the market price for allowances will be equal to the marginal abatement cost, which will be constant across emitters and across time (in discounted terms). In our laboratory setting, this price is $10.25. The variability of price away from the efficient price induces some inefficient abatement behavior and also increases price risk for market participants. Another possible effect of excessive price variability is that it may induce speculation on the value of allowances, based on the expectation that the quantity MSR will induce a rising pattern in allowance prices. So, for example, in an environment of excess allowances, the quantity MSR may be expected to lower auction quantities in subsequent periods. This adds a speculative value to holding allowances that could drive the price farther away from the efficient price. Other things equal, lower variability of price around the optimal price will improve market performance.

Result 3 (Price Variability): Observed auction prices exhibit significantly less variability with a price collar than is the case with a quantity collar or with no collar.

Support: Table 6 reports the results for the test of differences in price variability between treatments. The permutation test results make plain that the price collar performs significantly better than both the quantity MSR and the no MSR policy regimes in reducing price volatility. Figure 4 shows how participants in the price collar sessions are able to use allowance banking to smooth prices toward the optimal price relative to either the myopic (no banking) equilibrium or the quantity collar mechanism. What is more, this lower price variability does not come at the expense of effective price discovery or efficiency. Price variability is less, and the price stays closer to the optimum, even though the price is generally not constrained by the price collar itself.

Table 6. Stratified Permutation Test of Differences in the Mean Squared Difference of Price from the Optimal Price

Comparison	SD Price – SD Qty.	SD Price – SD No Policy	SD No Policy – SD Qty.
Difference(p-value)	-12.6 (0.001***)	-9.83 (0.003***)	-2.81 (0.506)

Significance level: * - 10%, ** - 5%, *** - 1%.

VI. Discussion and Extensions

The laboratory results offer several insights into the performance of alternative price-based and inventory-based adjustments to auction quantities that respond automatically to measures of market tightness. The experiments are run in an environment with an initial oversupply of allowances under a cap that is declining gradually over time. In a market where agents have long foresight and no constraints on intertemporal smoothing, we would expect to see a price very close to the long-run marginal abatement cost in every period. Our laboratory sessions with no policy intervention do show a pattern of agents smoothing short-run variations in permit prices relative to the no-banking (completely myopic) equilibrium. The subjects in the no-policy sessions did not exhibit sufficiently long-run foresight to keep prices stable throughout the experimental session. Series A sessions show a long-run downward trend in prices, while series B sessions had a generally increasing trend on average. It is of considerable interest that the implementation of a price collar induced more effective smoothing in both the long- and short-run horizons even though the collar was rarely binding. This suggests that the price collar acts through some other mechanism than through binding the permit prices to the specified range. The subjects may see the collar as reducing future price risk, as providing a value signal, or as a signal of policy commitment.

A second notable result is that quantity-based MSR may not be an effective way to reduce the *bank* of unused allowances. In fact, the effect that the policy has on banking behavior depends critically on particular characteristics of the allowance market. In series A, where the no-policy price is declining over time, the quantity MSR raises permit prices but results in a private permit bank that is nearly identical to the private bank under the price collar. In series B, where no-collar auction prices rise on average during the session, the quantity MSR has a relatively dramatic effect in lowering the private bank relative to the other treatments.

The justifications for the quantity MSR appear to rest on the assumption that the primary effect of the mechanism will be to force agents to use their banks rather than to buy allowances at auction, as fewer will be available. While there is some evidence of this in our sessions, there is more going on. In our sessions, as the quantity MSR reduces auction quantities, it raises the

market price above even the no-banking (perfectly myopic) equilibrium. This increases the opportunity cost of using permits, due either to the speculative value of a permit bank or to the subjects seeing the longer-term consequences of the policy. This causes the subjects to reduce their levels of production relative to efficient production levels, especially during periods when output prices are high. Reduced production during high value periods imposes significant efficiency penalties. Furthermore, once the bank falls below the low bank trigger point, prices fall well below the dynamically optimal price.

Another issue in the design of the quantity MSR involves the timing of the cycling of allowances into and out of the regulator's reserve. Although the quantity MSR could theoretically work to ease a shortage as well as reduce a current large bank, justification given for the quantity MSR, the existing surplus of issued EU ETS allowances over current compliance requirements has led to a relatively large bank of allowances and concerns over the price being too low. Our sessions raise some concern over how the quantity MSR will behave as the end of the current policy horizon approaches. If the reserve is exhausted, which would happen only under low private bank conditions, then allowance prices would be subject to considerable upside risk, since there are no other mechanisms for ameliorating the short-run fixed supply of allowances.[25] If, on the other hand, the reserve is large as the end of the regulatory period approaches, then allowance prices are subject to a collapse similar to the end of Phase I of the ETS unless participants expect the new policy regime to maintain the value of allowances by tightening the cap (i.e., retiring some portion of the reserve) or to carry the reserve indefinitely into the future. Given that it is not possible to time the reserve stock so that it will be close to zero at the end of the regulatory horizon, market behavior will reflect expectations about post-horizon treatment of allowances well before the program's final year. Considerable regulatory risk will be built into allowance prices.

To summarize, in some sessions, the delayed auction quantity reductions triggered by the quantity collar seem to have induced a "panic" in a tight market instead of merely tightening up a loose market with high initial endowments. This effect of the proposed policy merits further investigation with different experiment conditions and should not be dismissed as resulting from irrational reactions of inexperienced traders.

---

[25] This result is consistent with both the theoretical results in Salant (2015) but also with the stochastic optimization modeling experiment in Fell (2015).

The quantity collar mechanism under consideration by the EC is clearly a response to concerns that the current EU ETS allowance price is too low (European Commission, 2014a, 2014b). The policy discussions over low allowance prices place a strong weight on policies that support the price only indirectly through quantity adjustments. While it is well understood that the long-run aggregate supply of allowances will directly affect the price of allowances, it is less clear that short-run price adjustments can be effectively implemented in this way, at least not without the danger of substantial unintended consequences. Market participants are forward looking, if imperfectly so, and our experiments, along with other theoretical and modeling results, seem to show that there are a number of circumstances where efforts to manage the price through short-run shifting of allowance liquidity may run into difficulty by pushing against incentives to smooth costs over time. Our results confirm that a price collar mechanism may be less prone to these unintended consequences and may be a more robust way to ensure that the allowance price path supports long-run cost minimization.

Since both price and quantity collars are implemented with a market stability reserve of allowances held by the regulator, it would be useful to consider a hybrid policy that is based on a common MSR. Subsequent experiments could be conducted to determine whether the performance of a quantity collar might be improved with the incorporation of auction quantity adjustments that are triggered by auction clearing prices or sharp changes in those prices that would occur in the absence of further in-auction adjustments.[26]

Finally, the experiments reported here do not apply directly to the EU's currently implemented "backloading" policy of removing large blocks of allowances from upcoming auctions and then loading the reserve quantities into auctions held a few years later.[27] The relevant EU staff document recognizes that such a shift should have no effect on permit prices "in a perfect market" with foresight, but EC staff computer simulations generate substantial short-term price effects using various expectations assumptions (European Commission, 2014b). We are planning a follow-up experiment to evaluate the effectiveness of backloading as a way of correcting an initial oversupply of allowances. A second treatment with some uncertainty about the timing of allowance resupplies from the regulator's backloaded stock could also address some issues associated with regulatory risk.

---

[26] Computer simulations of such a hybrid policy have produced promising results, as reported at the DIW conference on MSR policies in Berlin, September 2014.

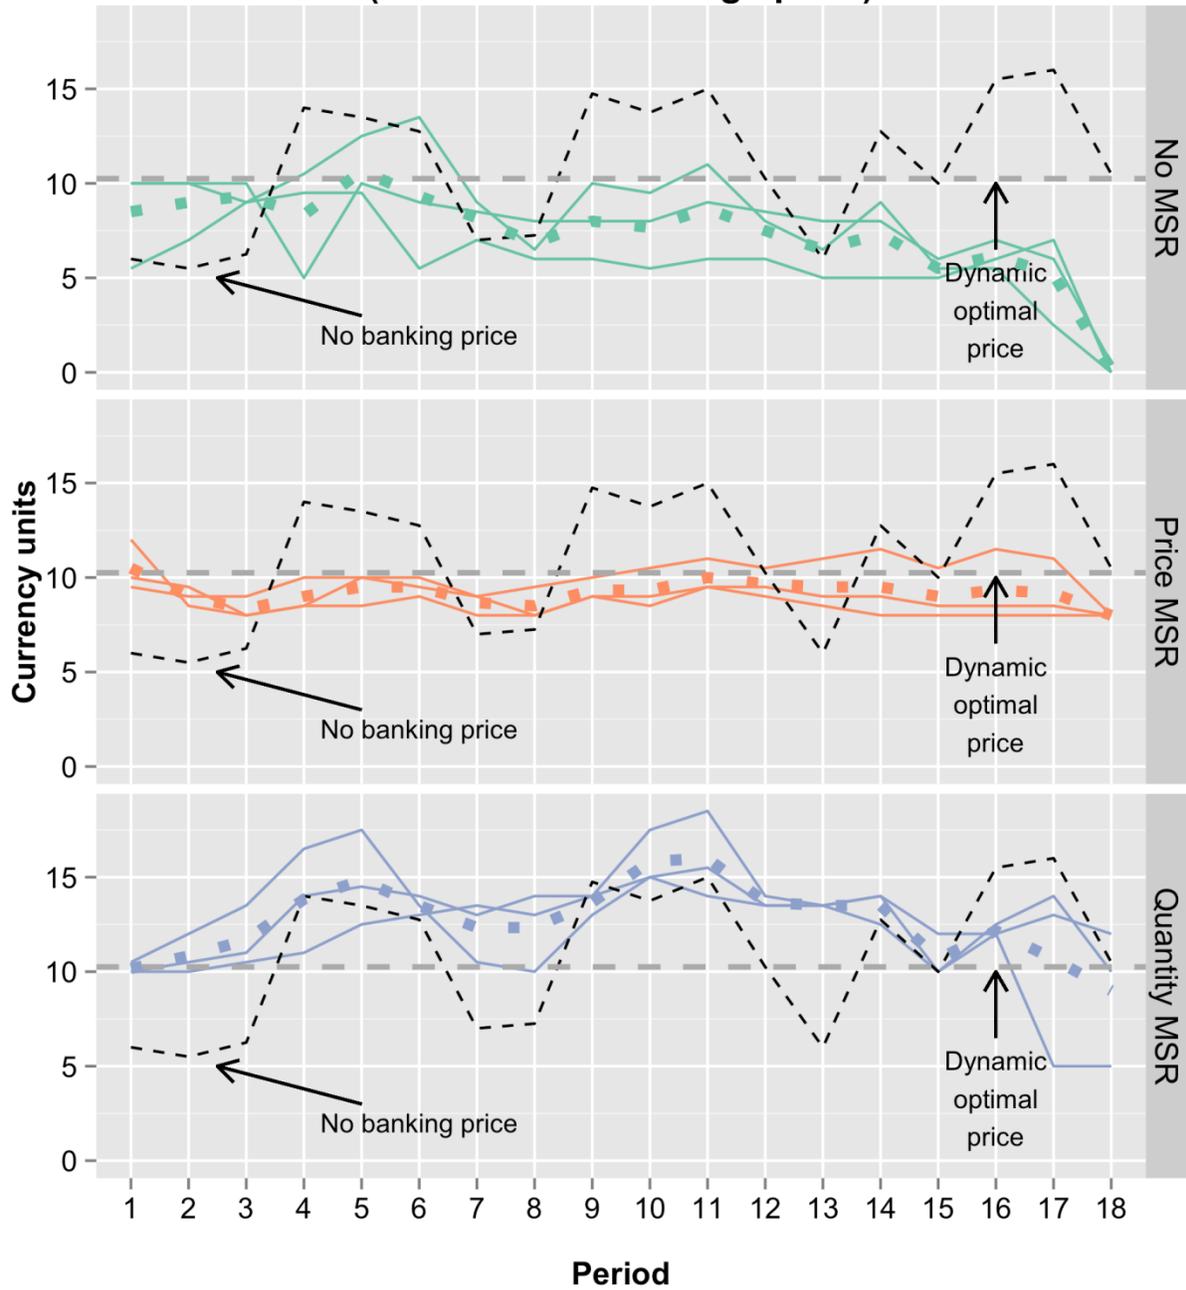[27] Commission Regulation (EU) No. 176/2014 of 25 February 2014.

References

Armantier, O., C.A. Holt, and C. Plott. 2013. A Procurement Auction for Toxic Assets with Asymmetric Information, *American Economic Journal: Microeconomics*, 5(4), 142–162.

Ausubel, L.M., and P. Milgrom. 2005. The Lovely but Lonely Vickrey Auction, in P. Cramton, Y. Shoham, and R. Steinberg, eds., *Combinatorial Auctions*, Chapter 1. Cambridge, MA: MIT Press.

Betz, R., J. Cludius, and A. Schopp. 2015. *The Role of the Financial Sector in European Union Emission Trading*. Unpublished manuscript.

Bodsky, R., D. Donato, K. James, and D. Porter. 2012. *Experimental Evidence on the Properties of California's Cap and Trade Price Containment Reserve*. Orange, CA: Economic Science Institute, Chapman University.

Burtraw, D., J. Goeree, C.A. 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.

Burtraw, D., A. Löfgren, and L. Zetterberg. 2014. *A Price Floor Solution to the Allowance Surplus in the EU Emissions Trading System*, Brief 14-02. Washington, DC: Resources for the Future.

Burtraw, D., K. Palmer, and D. Kahn. 2010. A Symmetric Safety Valve, *Energy Policy*, 38(9): 4921–4932.

Cason, T.N. 1995. An Experimental Investigation of the Seller Incentives in the EPA's Emission Trading Auction, *American Economic Review*, 85(4), 905–922.

———. 2010. What Can Laboratory Experiments Teach Us about Emissions Permit Market Design? *Agricultural and Resource Economics Review*, 39(2), 151–161.

Cummings, R., C.A. Holt, and S.K. Laury. 2004. Using Laboratory Experiments for Policy Making: An Example from the Georgia Irrigation Reduction Auction, *Journal of Policy Analysis and Management,* 3(2), 341–363.

Davis, D.D. and C.A. Holt. 1993. *Experimental Economics*. Princeton NJ: Princeton University Press.

European Commission. 2014a. *Commission Staff Working Document: Impact Assessment.* SWD(2014) 17 final, Brussels.

———. 2014b. *Commission Staff Working Document: Impact Assessment*. SWD(2014) 50 final, Brussels.

———. 2012. *Report from the Commission to the European Parliament and the Council: The State of the European Carbon Market in 2012*. COM(2012) 652 final, Brussels.

Fankhauser, S., and C. Hepburn. 2010. Designing Carbon Markets, Part I: Carbon Markets in Time, *Energy Policy*, 38(8), 4363–4370

Fell, H. 2015. Comparing Policies to Confront Permit Over-allocation, Discussion paper 15-17. Washington, DC: Resources for the Future.

Fell, H., D. Burtraw, R. Morgenstern, and K. Palmer. 2012. Soft and Hard Price Collars in a Cap-and-Trade System: A Comparative Analysis, *Journal of Environmental Economics and Management*, 64(2), 183–198.

Fell, H., and R. Morgenstern. 2010. Alternative Approaches to Cost Containment in a Cap-and-Trade System, *Environmental and Resource Economics*, 47 (2), 275–297.

Friesen, L., and L. Gangadharan. 2013. Environmental Markets: What Do We Learn from the Lab? *Journal of Economic Surveys*, 27(3), 515–535.

Goeree, J.K., and C.A. Holt. 2010. Hierarchical Package Bidding: A Paper & Pencil Combinatorial Auction, *Games and Economic Behavior*, 70(1), 146–169.

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*, Report Prepared for the New York State Energy Research and Development Authority, Albany.

Lopomo, G., L. Marx, D. McAdams, and B. Murray. 2011 Carbon Allowance Auction Design: An Assessment of Options for the United States, *Review of Environmental Economics and Policy*, 5 (1), 25–43.

Neuhoff K., A. Schopp, R. Boyd, K. Stelmakh, and A. Vasa. 2012. *Banking of Surplus Emissions Allowances: Does the Volume Matter?* DIW Discussion Papers No. 1193. Berlin: German Institute for Economic Research.

Perkis, D.F., T.N. Cason, and W.E. Tyner. 2015. An Experimental Investigation of Hard and Soft Price Ceilings in Emissions Permit Markets, *Environmental and Resource Economics*, forthcoming.

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.

Roberts, M., and M. Spence. 1976. Effluent Charges and Licenses under Uncertainty, *Journal of
    Public Economics* 5, 193–208.

Rubin, J. 1996 A Model of Intertemporal Emission Trading, Banking, and Borrowing, *Journal of
    Environmental Economics and Management*, 31, 269–286.

Salant, S. 2015. What Ails the European Union's Emissions Trading System? Two Diagnoses
    Calling for Different Treatments, Discussion paper 15-30. Washington, DC: Resources
    for the Future.

Shobe, W., C. Holt, and T. Huetteman. 2014. Elements of Emission Market Design: An
    Experimental Analysis of California's Market for Greenhouse Gas Allowances, *Journal
    of Economic Behavior and Organization*, 107(Part A), 402-420.

Shobe, W., K. Palmer, E. Myers, C. Holt, J. Goeree, and D. Burtraw. 2010. An Experimental
    Analysis of Auctioning Emission Allowances Under a Loose Cap, *Agricultural and
    Resource Economics Review*, 39(2), 162–175.

Taschini, L., S. Kollenberg, and C. Duffy. 2014. *System Responsiveness and the European
    Union Emissions Trading Scheme*, Policy Paper. London: Centre for Climate Change
    Economic and Policy, Grantham Research Institute on Climate Change and the
    Environment.

Trotignon, R., F. Gonand, and C. de Perthuis. 2014. *EU ETS Reform in the 2030 Climate and
    Energy Package: First Lessons from the ZEPHYR Model*. Climate Economics Chair
    Policy Brief 2014-01. Paris: Dauphine University.

Appendix: Graphs of All Session Observations

Session Group A - Auction Prices
(Dots indicate average price)

Session Group B - Auction Prices
(Heavy line is average price)