

Mind your Ps and Qs! An Experiment on Variable Allowance Supply in the US Regional Greenhouse Gas Initiative

Lana Friesen[†], Lata Gangadharan[‡], Peyman Kheizr[†], and Ian A. MacKenzie^{*†}

[†]School of Economics, University of Queensland, Brisbane, Australia

[‡]Department of Economics, Monash University, Australia

March 31, 2020

Abstract

Using an experimental approach, we investigate the new institutional design for the US Regional Greenhouse Gas Initiative (RGGI). The proposed scheme incorporates two allowance reserves that adjust the initial supply of allowances in the event of unexpectedly high or low allowance demand. In particular, allowance supply is increased when the initial clearing price is above a pre-determined upper trigger price, and decreased when the initial clearing price is below a pre-determined lower trigger price. We provide evidence that these two trigger prices act as focal points: the distribution of clearing prices is bimodal and aligns with the trigger prices. We also show that decreasing the range between the two trigger prices increases total revenue but decreases allocative efficiency. Importantly, we find the regulation is more sensitive to changes in trigger prices than reserve quantities.

Keywords: supply reserve; pollution allowances; experiment.

JEL Classification: C91; C92; Q58

*Corresponding author; e-mail: i.mackenzie@uq.edu.au. Phone: +61 7 336 56619.

1 Introduction

A common feature within existing pollution markets is the ability of a regulator to adjust allowance supply. These supply policies are typically incorporated within the initial allocation mechanism of each market: normally, pollution auction rules are adjusted so that the allowance supply can vary in the face of unexpected allowance demand. An increase in supply may assist in reducing firms' abatement costs whereas a reduction in supply may incentivize innovation in pollution abatement. Such adjustable allowance policies are used in the US Regional Greenhouse Gas Initiative (RGGI), European Union Emissions Trading Scheme (EU-ETS), and the California Cap-and-trade Program (Shobe et al., 2014; Perino and Willner, 2016; Burtraw et al., 2018).¹

In this article, we investigate one prominent design that has recently been adopted by the US Regional Greenhouse Gas Initiative (RGGI): the use of dual *allowance reserves* “Qs” and *trigger prices* “Ps”. An allowance reserve is a separate quantity of allowances that can be used to adjust the initial supply either up or down (Murray et al., 2009; Burtraw et al., 2018). Consequently, a reserve will be used to increase supply when the initial allowance price is sufficiently high: the initial price has to be larger than a ‘trigger’ price, which has been predetermined by the regulator. An allowance reserve can also be used to reduce supply, which works in a similar but opposite manner: if the initial allowance price is too low—below a set trigger price—the reserve will withdraw allowances from the supply.² We find that the regulator's choice of “Ps and Qs” has a pivotal effect on the regulatory outcome. In particular, we find the regulator's choice of trigger prices “Ps” act as focal points for the firms, which determines the allowance clearing price and thus influences revenue generation and allocative efficiency.

The actual RGGI design, proposed for the beginning of 2021, has been developed to not only contain firms' costs (by maintaining a pre-existing upper allowance reserve) but also to

¹In the EU-ETS, the Market Stability Reserve (MSR) adjusts allowance supply when the number of allowances in circulation falls outside a pre-defined range. The California Cap-and-trade Program incorporates the Allowance Price Containment Reserve (APCR) that allows sales of allowances given a sufficiently large clearing price in the preceding auction. See https://ec.europa.eu/clima/policies/ets/reform_en and <https://ww3.arb.ca.gov/cc/capandtrade/reservesale/reservesale.htm>.

²A reserve will increase supply until either (i) the price drops to the level of the trigger price or (ii) until the reserve is fully exhausted. Equally, a reserve will decrease supply until either (i) the clearing price equals the trigger price or (ii) the allowance reserve reaches its maximum capacity. Hence, if demand is sufficiently high or low, it is possible for the final clearing price to be outside the range of the trigger prices. Therefore, this design involves a “soft” price floor or ceiling. In contrast, a hard floor or ceiling strictly controls the price to be within that range.

contain emissions (by introducing a lower allowance reserve). This is observed within Figure 1.




Figure 1: The RGGI 2021 Quarterly Auction Supply. The projected base cap amount is approximately 15 million allowances with a Cost Containment Reserve (CCR) of 7,514,778 allowances and an Emissions Containment Reserve (ECR) of 6,845,333 allowances. Trigger prices for the release of the CCR is \$13 per ton while use of the ECR begins at \$6 per ton, and auction reserve price of \$2.38 per ton. Source: RGGI (2017)

RGGI’s initial allowance supply is projected to be around 13-15 million allowances per auction (where each allowance accounts for one ton of CO₂).³ The upper reserve denoted by the Cost Containment Reserve (CCR) holds a stock of 7,514,778 allowances that can be released to the auction if the initial clearing price is higher than the trigger price of \$13 per allowance. The lower reserve—denoted by the Emissions Containment Reserve (ECR)—has a capacity of 6,845,333 allowances that can be withdrawn from the market if the initial clearing price is below \$6 per allowance. Although this novel regulatory design will commence in 2021, it is *a priori* unclear how the design of multiple allowance reserves—with corresponding trigger prices—will affect the functioning of the auction and the regulatory scheme in general. In particular, it is unknown how the interaction and combination of these two allowance reserves will affect the clearing prices, revenue, and efficiency. Our paper, then, is the first analysis of this new dual allowance reserve institution.

In this article, we develop an experimental analysis of variable allowance supply, which

³The total annual cap for 2021 is 75,147,784 allowances, to be divided among the quarterly auctions, although this will be adjusted to take into account banked allowances already within the system. For details see RGGI (2017).

underpins the core institutional design of RGGI. In our stylized framework, we examine how the design of both an upper and lower allowance reserve (with associated trigger prices) alters the auction outcome. Our aim within this article, then, is to investigate how the regulator's design choices—of trigger prices, initial supply, and allowance reserves—impact the auction outcome. We find evidence that the regulator's chosen trigger prices act as focal points: the distribution of clearing prices is bimodal, which aligns with the trigger prices. We provide evidence that this is both an institutional design feature of having two allowance reserves and a consequence of buyers' bidding strategies within the regulatory system. Intuitively, if firms increase their bids in order to increase the clearing price up to the trigger price then they either (i) obtain more allowances at the lowest cost (the CCR is activated) or (ii) the threat of allowances being withdrawn is removed (the ECR is not activated). Furthermore, the equilibrium clearing prices are positively related to independent changes in both trigger prices.

We also find the range between the two trigger prices—what we denote as the *trigger-price collar*—has significant implications for the outcome of the scheme. As the width of the trigger-price collar decreases, i.e., the “soft” price floor and ceiling move closer together, we find this increases total revenue and the amount of allowances sold, but decreases allocative efficiency. Throughout our analysis, we find the scheme is more sensitive to changes in trigger prices than changes in quantities. Holding the maximum number of allowances constant, we find evidence that increases in the initial cap (with a smaller CCR and constant ECR) decreases the final clearing price as well as increasing the number of equilibrium allowances sold. However, we find that changes in quantities have *no significant effect* on revenue or allocative efficiency.

As our evidence highlights the focal-point characteristics of the regulator's trigger prices, we complement the prominent literature on the so-called *focal-point hypothesis* (e.g., Isaac and Plott, 1981; Coursey and Smith, 1983; Scherer and Ross, 1990; Knittel and Stango, 2003; Gode and Sunder, 2004; Engelmann and Müller, 2011; Rey and Tirole, 2019). This hypothesis argues non-binding price controls—usually “hard” price ceilings—may create a focal point in pricing decisions and thereby reduce competition. Although empirical evidence does exist for focal points in a limited number of markets (e.g., Knittel and Stango, 2003), the experimental laboratory evidence generally finds no collusive focal-point effects (Isaac and Plott, 1981; Coursey and Smith, 1983; Engelmann and Müller, 2011). In contrast to this literature, we find evidence of focal-point effects within a trigger-price system. This not only strengthens the argument for the existence of focal points but also widens the argument of the focal-point hypothesis from traditional non-binding “hard” price ceilings to incorporate the use of trigger prices, which are essentially “soft” price floors and ceilings, where the clearing price can be above or below

the predetermined trigger prices.

Our main innovation is to be the first analysis of dual allowance reserves. In doing so, we are able to analyze how the regulator's institutional design parameters—both separately and in combination—affect the operation of the regulatory scheme. More generally, our analysis also advances the literature on pollution market supply-side policies. The original idea of combining both prices and quantities to regulate can be traced back to Roberts and Spence (1976). Since then a broad literature has developed to investigate how alternative supply-side policies perform with respect to cost effectiveness. Murray et al. (2009) provided the first structured arguments for the use of allowance reserves and a number of studies have consequently investigated how a *single* allowance reserve—to ensure cost containment—would perform (Burtraw et al., 2010; Fell and Morgenstern, 2010; Grull and Taschini, 2011; Fell et al., 2012; Shobe et al., 2014; Stranlund et al., 2014; Kollenberg and Taschini, 2016; Khezzr and MacKenzie, 2018; Salant et al., 2019). Overall, this literature tends to find that adjusting allowance supply assists in cost containment and can improve welfare.⁴ Yet in our analysis, with dual allowance reserves, we find this may not always hold: the focal-point effects could mean that firms' costs increase if the regulator decides to raise *either* of the trigger prices. In terms of methodologies, the majority of this literature uses either theoretical or simulation modeling because obtaining data for these policies is extremely challenging: individual bidding behavior and actions are private information and the counterfactual is not known, which makes it hard to obtain a meaningful comparison. Consequently, there is a branch of literature that is now focusing on the experimental evidence of supply-side policies. The major benefit is that this can provide evidence of agent behavior within a controlled environment in order to test the impacts of regulatory design (Shobe et al., 2010, 2014; Holt and Shobe, 2016; Perkis et al., 2016; Burtraw et al., 2018). To date, however, the analysis of dual allowance reserves has been neglected.⁵

⁴A literature also exists that focuses on endogenous supply in auctions, where supply is chosen after bids have been submitted (Back and Zender, 2001; McAdams, 2007; Damianov and Becker, 2010) or allowed to weakly increase (LiCalzi and Pavan, 2005). This literature finds that allowing for forms of endogenous auction supply helps minimize demand reduction within the auction and thereby increasing the clearing price. Note, in contrast, that the allowance reserves adopted within the RGGI framework are very much distinct from what has been analyzed thus far: the use of allowance reserves—with fixed capacities and trigger prices that are known *ex ante* by buyers—create specific bidding strategies that have been previously unidentified. In particular, we provide evidence that these allowance reserves promote focal points in clearing prices, something that is not observed in alternative supply systems.

⁵Although Burtraw et al. (2018) focused on analyzing three ECR treatments relative to a no-ECR policy, they do not provide an investigation of dual allowance reserves nor the ramifications of the regulator's institutional design parameters. Stranlund et al. (2014) considers price controls and banking within an allowance market but their price controls are “hard” price floors and ceilings, which allow unlimited changes within allowance supply.

Our contribution is therefore in providing a greater understanding of the role and ramifications of allowance reserve policies and the associated focal points that are created. We find that the regulator’s pre-determined choice of initial allowance supply, trigger prices, and allowance reserves, have important implications for the functioning of the regulatory scheme. The findings will provide further knowledge for the designers of prominent allowance markets. Our article is organized as follows. In Subsection 1.1, we provide a background on the institutional design within RGGI. In Section 2 we outline the experimental design. Sections 3 and 4 present the results, while Section 5 provides some concluding remarks.

1.1 Background: Allowance Allocation in the Regional Greenhouse Gas Initiative (RGGI)

The US Regional Greenhouse Gas Initiative (RGGI) is a regional cap-and-trade market that commenced in 2008 and regulates emissions from the utility sector in the Northeastern US states. The majority of allowances within RGGI are allocated via a quarterly auction, where, on average, 43 firms participate in bidding (2017-2019 inclusive) and has generated a cumulative revenue of approximately US\$3.5 billion since 2008.⁶ The auction format used is the multi-unit uniform-price auction, which is the predominant method to distribute pollution allowances (Lopomo et al., 2011). Each individual auction generates around US\$75 million in auction revenue, which is invested in mainly energy and consumer programs.

As the program has developed over time—with two major program reviews in 2012 and 2017—the auction rules have also developed. Unlike conventional allowance auction designs that usually offer a fixed amount of allowances at each auction, the auction rules of RGGI have been developed to vary the supply of allowances. Within the first program review in 2012 (and implemented in 2014), the Cost Containment Reserve (CCR) was established to reduce significant price increases (RGGI, 2013). The CCR was originally a fixed amount of 10 million allowances with the trigger price being initially set at US\$4 per allowance with increases year-on-year (the initial allowance supply was also continually decreasing year-on-year). Within the 2017 Program Review (RGGI, 2017), a similar but opposite reserve was developed. From 2021, RGGI will implement the Emissions Containment Reserve (ECR). The rationale for establishing an ECR is to reduce supply of allowances at sufficiently low prices, which may therefore ensure a sustained incentive to invest in abatement technology. The CCR and ECR have annual capacities—10% of the regional (non-adjusted) cap—so if the allowances are exhausted (filled to capacity) in a previous auction then these reserves are no

⁶See <https://www.rggi.org/auctions/auction-results>.

longer available for use until the next calendar year.⁷

As described earlier, and shown in Figure 1, the RGGI design in 2021 is fairly symmetric in terms of the size of allowance reserves with both the ECR and CCR containing approximately half of the number of allowances contained in the initial supply. The trigger prices for the ECR and CCR will initially be set at \$6 and \$13 per allowance, respectively, and will increase by 7% per annum. Yet, note although these trigger prices increase at a constant rate each year, the absolute size of the trigger-price collar increases over time due to different baselines. This is shown in Figure 2.




Figure 2: The RGGI Trigger Price Levels. The CCR trigger prices (black line) commenced in 2014 and the ECR trigger prices (red dashed line) will commence in 2021. The absolute size of the trigger-price collar is increasing over time.

Given our evidence within this article—along with the institutional development of RGGI presented in Figure 2—we find that as the trigger-price collar increases, we would expect, *ceteris paribus*, to see increases in allowance prices. Furthermore, if the width of the trigger-price collar increases, *ceteris paribus*, we expect that there may be an increased likelihood of a reduction in allowances sold but an increase in allocative efficiency.

Our stylized control treatment resembles RGGI in that we have symmetric levels for the ECR and CCR. It is clear from our analysis that how the trigger prices are chosen—both separately and collectively—have important implications for revenue and allocative efficiency. What our evidence suggests is that, within RGGI, the regulator could make improvements in its criteria of importance (such as revenue generation, efficiency, etc) by adjusting the trigger prices in both an absolute and relative manner. However there may exist tradeoffs. A tighter

⁷Note the stock of the ECR is 10% of the regional (non-adjusted) cap but excluding the budgets of Maine and New Hampshire as they do not intend to participate in the ECR process.

trigger-price collar, may increase revenue but at the expense of reduced allocative efficiency.

2 Dual Allowance Reserves

In this section, we develop a stylized framework that reflects the key institutional design features proposed within RGGI. This framework then informs our experimental design.

2.1 Preliminaries and Notation

Consider a regulator that has the ability to sell a maximum of $\bar{Q}$ allowances within a single auction to a set of potential firms $\{1, 2, \dots, n\}$. Firm i has a private value v^i for the allowances and we assume these values are distributed independently from a known distribution function.⁸ We interpret these values as firms' (avoided) heterogeneous marginal abatement costs.⁹ We abstract from analyzing the secondary market as the objective of the article is to focus on the use of allowance reserves within the initial allocation process, as proposed by the Regional Greenhouse Gas Initiative (RGGI). Note, however, that if a secondary market were to be included after the auction, then—through the process of backward induction—each firm's value of an allowance in the auction could be interpreted as their expected value of an allowance on the secondary market. Consequently, the fundamental elements of the auction do not change (with only a change in the interpretation of firms' values). As is standard within allowance market auctions, we investigate a sealed-bid uniform-price auction. Once firms submit their sealed bid-schedules, the regulator orders firms' bids from the highest to the lowest and the winning firms receive a share of the allowances at a price equal to the highest losing bid. If there are ties, then the regulator randomly chooses the winners. We normalize the reserve

⁸The alternative would be to assume a common values setup, yet this is less realistic in the context of pollution abatement because this would restrict all firms to an identical marginal abatement cost.

⁹Note that v^i does *not* mean that firms' marginal abatement cost curves are globally horizontal. First, if the marginal abatement cost curve is continuous and convex then within the static auction their value is simply their current marginal cost (i.e., no changes in abatement occur when the auction is running). The value is therefore determined by their historical abatement and purchases of allowances prior to the auction commencing: we model this by allowing the value to be randomly redrawn at the beginning of each auction round. Second, this value is also consistent with the standard construction of abatement curves (e.g., McKinsey and Company, 2016), where a firm's long-run abatement cost curve is globally convex but composed of a step function that discontinuously jumps when new, more expensive, technologies are adopted. The value can therefore be interpreted as the forgone cost of using a specific abatement technology up to capacity. As our focus is on subject behavior in a single auction, we abstract from the banking of allowances, but note that firms' values are general enough to be interpreted as their expected value of an allowance, where banking occurs in a later stage. See Stranlund et al. (2014) for a discussion on banking in an experimental setting.

price to zero.

We allow the regulator to initially offer a volume of allowances denoted by Q_2 , which can be either increased up to $\bar{Q}$ —if the clearing price is sufficiently high (using the Cost Containment Reserve)—or decreased to Q_1 —if the initial clearing price is sufficiently low (using the Emissions Containment Reserve). This is shown within Figure 3.




Figure 3: Variable Allowance Supply. In (a), the important parameters of the trigger prices $\{\tilde{p}_1, \tilde{p}_2\}$ and the quantities $\{Q_1, Q_2\}$, which determine the initial offering of allowance in the auction as well as the size of the ECR and CCR for a fixed and given $\bar{Q}$. The auction reserve price is normalized to zero. In (b), if demand occurs as expected ($D_E(p)$), allowance supply is given by Q_2 . When demand is higher than expected ($D_H(p)$) allowance supply increases from Q_2 to Q_H . When demand is lower than expected ($D_L(p)$) allowance supply decreases from Q_2 to Q_L .

Formally, denote the size of the Cost Containment Reserve as $CCR \equiv |\bar{Q} - Q_2|$ and the Emissions Containment Reserve as $ECR \equiv |Q_2 - Q_1|$. The CCR becomes operational when the initial auction clearing price is above the CCR trigger price, which is denoted by $\tilde{p}_2$. Analogously, the ECR becomes operational if the initial auction clearing price is below the ECR trigger price, which is denoted by $\tilde{p}_1$.

If the initial clearing price in the auction is above the upper bound trigger price $\tilde{p}_2$, then the auctioneer releases a maximum of $\bar{Q} - Q_2$ allowances into the auction supply until either (i) the CCR is exhausted or (ii) until the clearing price decreases to equal the trigger price $\tilde{p}_2$, (whichever is first). If, on the other hand, the initial clearing price is sufficiently low, that is, below the lower bound trigger price $\tilde{p}_1$, a maximum of the ECR can be withdrawn from the auction until either (i) the ECR is at maximum capacity or (ii) the clearing price increases to

equal the trigger price $\tilde{p}_1$, (whichever is first). Figure 3 (b) shows what happens in different demand scenarios. If allowance demand is as expected ($D_E(p)$), then the regulator continues to supply Q_2 allowances. If demand is sufficiently high, ($D_H(p)$), then the initial price will be above the trigger price $\tilde{p}_2$, so supply increases. In this case supply will increase to point Q_H (rather than the full reserve being exhausted). If, on the other hand, demand is lower than expected ($D_L(p)$), then the initial allowance price is lower than the trigger price $\tilde{p}_1$, and allowances are withdrawn from supply. In this example supply falls to Q_L (and the reserve is not fully used).

Note that this supply structure generates “soft” price floors and ceilings. From Figure 3 (b), if $D_H(p)$ continually moved in the Northeast direction, it is clear to see that the reserve would eventually be fully exhausted and, consequently, the clearing price would increase above $\tilde{p}_2$. Analogously, if $D_L(p)$ moved in the Southwest direction, then, as some level, the the ECR would reach capacity and the price would decrease below $\tilde{p}_1$.

Comparison of Figures 1 and 3 shows the key features of RGGI’s variable allowance supply approach within our stylized framework; namely, allowing for two allowance reserves with distinct trigger prices. Consequently, the aim of this article is to investigate how the regulator’s chosen institutional parameters $\{\tilde{p}_1, \tilde{p}_2, Q_1, Q_2\}$ affect the regulatory outcome (in terms of final clearing price, revenue generation, allowances sold, and efficiency).

2.2 Experimental Design

To explore how the design of the RGGI regulations affect outcomes, we designed a stylized laboratory experiment that matches the key features proposed in the policy. In particular, all of the treatments involve a sealed-bid multiple-unit uniform-price auction with two trigger prices and an ECR and CCR as described earlier. Each auction involves four bidders in fixed groups bidding for a maximum of six allowances, where each bidder can demand at most two allowances with values randomly drawn from $U[0, 10.00]$ at the beginning of each round. The values are private, the same for both units of the allowance, and are randomly redrawn at the beginning of each round. We use a total of four bidders in each auction as this provides us with enough manageable buyers to replicate realistic interactions often observed in multi-unit auctions.¹⁰ Recall that within RGGI, the average number of active bidders is relatively low at

¹⁰ Allowance auction experiments have typically involved 4-10 subjects in a market, with between 2-4 distinct types (e.g., Shobe et al., 2010; Stranlund et al., 2011; Shobe et al., 2014; Stranlund et al., 2014; Holt and Shobe, 2016). More general multi-auction experiments with private values (e.g., Engelmann and Grimm, 2009; Pagnozzi and Saral, 2017), have studied only two bidders per market, considering this sufficient to capture the strategic elements. Both of these experiments also include flat demands for two units. Note that Engelbrecht-Wiggans et al.

43 (2017-2019): the lowest being 29 bidders in Auction 18 and, for the futures auction, there were only 4 bidders in Auction 10 (futures).¹¹

Table 1: Treatments and Parameters

Treatment label	Q_2	CCR	ECR	$\tilde{p}_2$	$\tilde{p}_1$	Number of groups (individuals)
Control	4	2	2	6.66	3.33	7 (28)
LowP2	4	2	2	5.00	3.33	7 (28)
LowP1	4	2	2	6.66	2.00	7 (28)
LargeQ2	5	1	2	6.66	3.33	7 (28)
SmallQ2	3	3	2	6.66	3.33	6 (24)

Our five treatments are summarized in Table 1. We design our baseline treatment, Control, such that the two trigger prices equally divide the value range, while the quantities are also equally and symmetrically divided. In two of our treatments, we independently lower the two trigger prices; LowP2 and LowP1. Since we only change one trigger price at a time, these two treatments are also equivalent to narrowing and widening the size of the trigger-price collar, respectively. Comparing Control with LowP2 and LowP1 allows us to examine if varying trigger prices, while keeping the structure of the allowance reserves the same, influences the auction outcome. We can also examine which of the two trigger prices would have a greater impact on outcomes.

In two other treatments, we keep the trigger prices the same as in the Control but vary the construction of the reserves to make them asymmetric. In particular, in the LargeQ2 treatment we increase the initial supply, while in the SmallQ2 treatment we decrease it. As we hold the maximum available allowances fixed, this means we have cases with a large initial supply, small CCR, and constant ECR (and vice versa).¹² A comparison of these two treatments with Control helps us understand the impact of changing the structure of the allowance reserve.

In each round bidders enter their demand schedule (or maximum willingness to pay) for the two allowances. They can bid any non-negative cent amount; i.e., we do not place any restriction on the bidding schedules, with no upper limit nor restriction that bids be the same

(2006) provide experimental field evidence that increasing the number of bidders does not eliminate the main bidding strategies (demand reduction) in multi-unit auctions. In our approach, subjects' values are randomly drawn, which provides a sufficient degree of heterogeneity and allows the showcasing of the main effects from these auctions.

¹¹See <https://www.rggi.org/Auctions/Auction-Results/Supply-Bid>.

¹²That is, in our design a larger (smaller) initial supply corresponds to a smaller (larger) CCR by definition, as the ECR and overall number of allowances is held constant.

for both units of the good. After all buyers have submitted their bids, the computer orders the eight bids in each group from highest to lowest. The quantity sold is determined according to the institutional features of each treatment as described below. In all cases, the auction clearing price is determined by the highest rejected bid but the quantity supplied may differ.

In all treatments, additional allowances are supplied if the interim price exceeds the trigger price $\tilde{p}_2$ until either the supply is exhausted or the clearing price falls to the trigger price. Equally if the interim price is below the trigger price $\tilde{p}_1$ supply is withdrawn until the price equals the trigger price or the reserve reaches capacity. In all treatments, the allowances are allocated to the highest bidders. In the case of excess demand, a random rationing rule is used which avoids “part” units of allowances being allocated as would occur under a proportional rule. After each round subjects are told the auction outcome, which is described by i) the market price, ii) the total number of allowances sold in the auction, and iii) the number of allowances they purchased in the auction. The bids and values of other bidders remain private information throughout. Bidders earn profits (or losses) for every allowance they purchase for a price less (more) than its value.

Note that while we use environmental terminology (e.g., allowances) in the article, the experiment itself is neutrally framed. Subjects are instructed that they are bidding for a non-descript good with a private value and the supply of the good is an increasing step function.¹³ All sessions were conducted at the University of Queensland in October 2019, using z-Tree (Fischbacher, 2007). Subjects were recruited from the general pool of subjects using ORSEE (Greiner, 2015) and were mostly undergraduate students, with about half of them studying business or economics degrees. We collected data from 136 subjects, each of whom participated in only one treatment. Table 1 reports the number of subjects and independent groups in each treatment. Around 40% of the subjects were male, and although most of the subjects had taken part in previous experiments, none had experience in auction experiments.

The experiment proceeds as follows. At the start of the session, detailed instructions are read aloud by an experimenter (always the same person) while the subjects follow along on their written copy. Subjects then answer a series of quiz questions (10-12 depending on the treatment), earning AU\$0.50 for each correct answer. The quiz questions check and reinforce understanding of the instructions. Next, subjects take part in the main auction task which is repeated over 25 rounds. Finally, to examine if attitudes towards risk influence individuals’

¹³Full experimental instructions for treatment Control are included in Appendix A. As can be observed, in order to ensure tractability and to avoid framing effects in the supply design, the instructions explained the regulatory design as a generic increasing stepped supply function. Subject understanding can be observed by noting that the modal allowances sold was Q_2 in all treatments, which is the initial allowance supply.

decisions in the auction, we elicit risk preferences using a scaled version of the Eckel and Grossman (2008) lottery-choice task, which is followed by a demographic questionnaire.

Subjects are paid their profits in five randomly selected auction rounds (chosen at the end of the entire experiment) plus their earnings from the risk task and quiz, in addition to a \$5 show up fee. Subjects earned AU\$28.09 on average (equivalent to US\$19 at the time of the experiment), and range from AU\$10.60 to 50.60. Each session lasted on average 90-120 minutes including payments.

3 Results: Market Level

In this section, we analyze results at the aggregate market level. We focus first on the fundamental auction outcomes of the allowance clearing price and the number of allowances sold. We then study two measures of auction performance: total revenue raised in the auction (i.e., price times the number of allowances sold), and allocative efficiency, which is defined as the percentage of the maximum possible surplus actually achieved in the auction. Summary statistics for these four key measures are reported in Table 2 disaggregated by treatment.

Our primary analysis relies on random effects panel regressions with standard errors clustered at the market level. These regressions include indicator variables for each treatment, except for the Control treatment, as well as the inverse of the period to account for learning. Results are reported in Table 3, where the coefficient estimates are interpreted as treatment effects *relative* to the Control treatment. The results show strong learning effects for all measures with the coefficient on *Inverse Period* strongly significant in all columns. In the text, we also report results from Mann-Whitney tests of treatment differences using each market (i.e. each independent group) as the unit of analysis. Unless stated otherwise, all p-values reported in the text are exact and from two-sided tests.

Table 2: Mean Value (Standard Deviation) of Key Measures

	Control	LowP2	LowP1	LargeQ2	SmallQ2
Clearing Price	4.454 (0.250)	4.133 (0.460)	3.757 (0.333)	4.015 (0.614)	5.076 (0.257)
Allowances Sold	3.903 (0.160)	4.280 (0.517)	4.006 (0.134)	4.503 (0.288)	3.180 (0.194)
Total Revenue	18.270 (1.899)	18.547 (4.037)	15.748 (1.667)	18.896 (4.101)	16.848 (1.729)
Allocative Efficiency (%)	95.389 (1.851)	94.219 (4.420)	97.437 (2.226)	96.639 (1.370)	95.138 (2.955)
Observations	7	7	7	7	6

Table 3: Main Treatment Effects: All Periods, Inverse Time Trend

	(1)	(2)	(3)	(4)
	Clearing Price	Allowances Sold	Total Revenue	Allocative Efficiency
LowP2	-0.321* (0.186)	0.377* (0.193)	0.278 (1.589)	-1.170 (1.707)
LowP1	-0.697*** (0.148)	0.103 (0.074)	-2.522*** (0.900)	2.048** (1.032)
LargeQ2	-0.439* (0.236)	0.600*** (0.117)	0.626 (1.610)	1.250 (0.820)
SmallQ2	0.622*** (0.132)	-0.723*** (0.093)	-1.422 (0.942)	-0.251 (1.301)
Inverse Period	-0.853*** (0.190)	-0.418*** (0.147)	-4.980*** (1.076)	-5.435*** (2.033)
Constant	4.584*** (0.096)	3.967*** (0.065)	19.030*** (0.713)	96.219*** (0.706)
Observations	850	850	850	850
Number of Groups	34	34	34	34
LowP2 = LowP1	0.063	0.149	0.072	0.068
LargeQ2 = SmallQ2	0.000	0.000	0.201	0.220

Results from random effects panel regressions with standard errors clustered at the market level.

The last two rows report p-values from Wald tests of treatment differences. The omitted base treatment is Control.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.1 Allowance Clearing Price

We begin by investigating how the implementation of two allowance reserves—with their associated trigger prices—alters the final clearing price.

Result 1. *In all treatments, the distribution of clearing price(s) is bimodal, and aligns with the relevant trigger price(s).*

Support. The distribution of final clearing prices is shown for each treatment in Figure 4, where the solid lines indicate the trigger prices in each treatment. In each treatment, this distribution is bimodal at the trigger prices in that treatment. Clearing prices above $\tilde{p}_2$ are rare, comprising less than 1% of all prices in each treatment except for LowP2. Even when $\tilde{p}_2$ is low, as in LowP2, the trigger price itself has considerable drawing power because even in this case only 5% of prices exceed this trigger price.

Result 1 suggests that a major implication of implementing dual allowance reserves is the clearing price will tend towards the closest trigger price: this occurs for either cost or emissions containment reserves. It is clear from observing Figure 4 that trigger prices may be a focal point/anchor in agents' bidding strategies as well as an institutional design feature (this is confirmed later in the article when we investigate individual bidder behavior). Intuitively, bidders have an incentive to try to bid the price up to the trigger price: this will release additional allowances at the lowest possible cost (for the CCR) or avoid a withdrawal of allowances (for the ECR) (Khezr and MacKenzie, 2018).

We now investigate what happens to the clearing price when the regulator's trigger prices are altered.

Result 2. *Lowering either $\tilde{p}_1$ or $\tilde{p}_2$ trigger prices significantly decreases the auction clearing price.*

Support. The average clearing price in Control is \$4.45, compared to \$4.13 in LowP2 and \$3.76 in LowP1. Using Mann-Whitney tests, the difference between Control and LowP1 is strongly significant ($p=0.0006$) but the difference between Control and LowP2 is not ($p=0.26$). The regression results reported in Column (1) of Table 3 show that lowering either trigger price significantly reduces the final clearing price, although the impact of lowering $\tilde{p}_1$ is stronger both in magnitude and in significance than lowering $\tilde{p}_2$.¹⁴

¹⁴Comparing the variance of clearing prices across treatments, we find that price variance is significantly lower in LowP2 compared to Control ($p = 0.0023$) but significantly higher in LowP1 versus Control ($p = 0.0973$). LowP1 and LowP2 are also significantly different ($p = 0.0006$). Price variance is not significantly different between Control and either LargeQ2 or SmallQ2 or between LargeQ2 and SmallQ2. This implies that the width of the trigger-price collar impacts the degree of price volatility, where a narrower collar reduces price volatility (and vice versa). On the other hand, changes in the quantities have no significant impact on price volatility.




Figure 4: Distribution of Prices

Result 2 shows that if either $\tilde{p}_1$ or $\tilde{p}_2$ decrease, then this will decrease the final clearing price. Given Result 1 shows that the modal clearing prices are the trigger prices, it is now intuitive that a decrease in either trigger price $\{\tilde{p}_1, \tilde{p}_2\}$ will generate lower clearing prices. In terms of the application to RGGI, we know the trigger prices for both reserves will be automatically increased by 7% per annum. Thus our evidence suggests, *ceteris paribus*, that as trigger prices increase there will be upward pressure on the final clearing price from *both* trigger prices. This is somewhat surprising. We would expect that $\tilde{p}_1$ may have a positive influence on the clearing price (as this has the ability to withdraw allowances), yet the trigger price for the CCR $\tilde{p}_2$ also has a positive influence on the final clearing price. Intuitively an incentive may exist to bid the price up to obtain additional allowances from the CCR.

Result 3. *A lower $\tilde{p}_1$ trigger price has a significantly greater dampening effect on prices than a lower $\tilde{p}_2$ trigger price, *ceteris paribus*.*

Support. The estimated treatment effect for LowP1 reported in Table 3 is more than double that for the LowP2 treatment. This difference is significant according to the p-value reported in the bottom panel of Table 3. It is worth remembering that for design reasons the size of the absolute reduction in LowP1 (1.33) is actually slightly smaller than in LowP2 (1.66).

The trigger price of the ECR has a significantly greater effect than the trigger price of the CCR. As the ECR trigger price is low it will be breached only for sufficiently low truthful bids or substantial bid shading. What this result suggests is that there is a stronger effect for increased bids in order to avoid allowance withdrawal (ECR) relative to the case where additional allowances are added due to the CCR. In relation to RGGI, if the regulator aims to have higher clearing prices over time, the biggest effect will come from increasing the ECR trigger price $\tilde{p}_1$ at a faster rate than the CCR trigger price $\tilde{p}_2$.

We now investigate the implications of adjusting Q_1 and Q_2 relative to the control treatment. Thus we will consider what happens when we change the initial allowance supply Q_2 and the lower allowance point Q_1 . Doing so, then, helps us capture the effects not only of changing the initial supply but also the relative sizes of the ECR and CCR. In particular, we hold the size of the ECR constant and allow for a “high” initial supply with a “low” CCR (LargeQ2) as well as a “low” initial supply with a “high” CCR (SmallQ2).

Result 4. *Increasing Q_1 and Q_2 (i.e., the initial supply is larger with a smaller CCR) significantly decreases the auction clearing price. Decreasing Q_1 and Q_2 (i.e., the initial supply is smaller with a larger CCR) significantly increases the auction clearing price.*

Support. The average clearing price in Control is \$4.45, compared to \$4.02 in LargeQ2 and \$5.08 in SmallQ2. Using Mann-Whitney tests, the difference between Control and SmallQ2 is

strongly significant ($p=0.0023$) while the difference between Control and LargeQ2 is borderline ($p=0.0973$). The regression results reported in Column (1) of Table 3 confirm these findings, with SmallQ2 significantly increasing the price relative to the Control, while LargeQ2 weakly decreases it. Intuitively there is less incentive to release the final allowance in the LargeQ2 treatment, relative to a stronger incentive to release the final *two* allowances in the Control treatment (i.e., greater incentive to hit the second trigger price in the latter), and there is more incentive to release the final *three* allowances in the SmallQ2 treatment, relative to the slightly weaker incentive of only releasing a final *two* allowances in the Control treatment (i.e., greater incentive to hit the second trigger price in the former).

Note that in RGGI, Q_2 is the intended supply for the auction. Thus as one would expect a lower (higher) intended supply will increase (decrease) the final clearing price. With a relative smaller CCR, there is less incentive to bid higher to release the smaller CCR, *ceteris paribus*. Consequently a smaller CCR with a larger intended auction cap will force the final clearing down.

3.2 Allowances Sold

We now want to investigate if there is any impact on the number of allowances sold in equilibrium if we introduce alternative allowance reserve architecture.

Result 5. *The modal number of allowances sold aligns with Q_2 in all treatments.*


Support. Figure 5 shows the distribution of the number of allowances sold in each treatment. The figure reveals a strong mode in all treatments at Q_2 . This suggests, as expected, that the placement of the intended supply Q_2 is an important determinant of the final number of allowances sold. Hence, if a goal is to *control the number of allowances sold*, the auctioneer should place greater focus on the choice of Q_2 , rather than the choice of Q_1 (*ceteris paribus*), i.e., the size of the ECR and CCR.

Let us now observe if the trigger prices have an impact on allowances sold.

Result 6. *Lowering the $\tilde{p}_2$ trigger price (i.e., reducing the price collar between $\tilde{p}_1$ and $\tilde{p}_2$) weakly increases the number of allowances sold. Lowering the $\tilde{p}_1$ trigger price (i.e., increasing the price collar between $\tilde{p}_1$ and $\tilde{p}_2$) has no significant effect on the number of allowances sold.*

Support. On average, 3.9 allowances are sold in the Control treatment, compared to 4.3 in LowP2, a weakly significant difference ($p=0.07$). In contrast, 4.0 allowances are sold in LowP1, which is nearly identical to the Control treatment and indeed the difference is not significant ($p=0.20$). The regressions reported in Column (2) of Table 3 support these findings

Distribution of Allowances Sold



Reference lines at relevant Q2 values

Figure 5: Distribution of Allowances Sold

with the estimated coefficient on LowP2 weakly significant, while the coefficient on LowP1 is not significantly different from zero.

The CCR in the relevant treatments, Control, LowP2 and LowP1, contains two allowances (allowances 5 and 6). This CCR is more likely to be released when the second trigger price is lower as shown in Figure 5. While the mode is the same across all three treatments ($Q=4$), there is a greater proportion of data at $Q = 5, 6$ in the LowP2 treatment. This is evidence that, by lowering the second trigger price, there is an increase in the likelihood that the second extra quantity will be released. For the second part of the result, lowering the first trigger price should have minimal impact on the number of allowances released. That is, the initial price is typically above the first trigger, whether it is \$2 or \$3.33. This is evidenced by the mean number of allowances sold in Table 2 (3.9 and 4.0 allowances in Control and LowP1 respectively).

Let us now investigate whether the regulator's choice of quantities determines the number of allowances sold.

Result 7. *Increasing (decreasing) Q_1 and Q_2 significantly increases (decreases) the number of allowances sold.*

Support. On average 3.9 allowances are sold in Control compared to 4.5 in LargeQ2 and 3.2 in SmallQ2. Both of these treatments are significantly different from Control and each other (all p -values < 0.002). These findings are further confirmed in column 2 of Table 3 where the coefficient on LargeQ2 is significantly positive relative to Control, while the coefficient on SmallQ2 is significantly negative. These results are highlighted in Figure 5, where as discussed earlier, the modal number of allowances sold aligns with Q_2 in all treatments.

This result shows that having a larger initial supply and lower CCR will result in an increase in allowances sold, and vice versa.

3.3 Total Revenue

Another goal of most regulators is to obtain sustainable levels of auction revenue. We now consider how alternative institutional designs impact total generated revenue. We first consider how changing the trigger prices affects total revenue.

Result 8. *Lowering the $\tilde{p}_1$ trigger price (i.e., increasing price collar between $\tilde{p}_1$ and $\tilde{p}_2$) significantly decreases total revenue. Lowering the $\tilde{p}_2$ trigger price (i.e., decreasing the space between $\tilde{p}_1$ and $\tilde{p}_2$) has no significant impact on total revenue.*

Support. Average total revenue in Control is 18.27 compared with 15.75 in LowP1, a significant difference ($p=0.01$). In contrast, average total revenue is 18.55 in LowP2, nearly

the same as in the Control treatment ($p=0.71$). The regression results in Column (3) of Table 3 confirm these findings with the coefficient on LowP1 significantly negative, while that for LowP2 is statistically zero.

The negative effect on revenue levels in LowP1 is expected given that clearing prices are significantly lower yet the number of allowances sold is statistically similar. For the second part of the result, the absence of any significant difference in revenue levels is not unexpected; while prices are significantly lower (albeit weakly so), the number of allowances sold is (weakly) significantly greater than in the Control treatment. That is, there are two offsetting forces at play when the second trigger price is lowered; more allowances are sold, but at a lower clearing price, leading to no significant change in revenue levels.

We now investigate whether the regulator's choice of Q_1 and Q_2 has implications for the total revenue generated.

Result 9. *Increasing or decreasing Q_1 and Q_2 has no significant impact on total revenue.*

Support. Average total revenue in LargeQ2 is 18.90, which is similar to the average of 18.27 in Control ($p=1.00$). While revenue in SmallQ2 is somewhat lower at 16.85, the difference is not significant ($p=0.18$). The results in Table 3 are consistent with this, with neither coefficient on LargeQ2 or SmallQ2 significant.

Again, this absence of any significant difference in revenue levels is not unexpected; in both treatments there are two offsetting forces at play—prices are significantly lower (higher) in the LargeQ2 (SmallQ2) treatment, but the number of allowances sold is significantly higher (lower). This result shows that the two effects exactly offset each other.

3.4 Allocative Efficiency

We now investigate how dual allowance reserves impact the allocative efficiency of the mechanism. We calculate allocative efficiency as follows (i) first, for each market, in each period, we calculate the sum of the private values of the bidders who successfully purchased allowances and (ii) second, for the same market and the same period, we calculate the sum of the *highest* private values of the *same number of bidders*. Allocative efficiency for a given market for a given period is then calculated as the ratio of (i) to (ii). Hence, allocative efficiency of 100% implies that the allowances were *actually* sold to the highest value bidders.

Result 10. *Lowering the $\tilde{p}_1$ trigger price (i.e., increasing the price collar between $\tilde{p}_1$ and $\tilde{p}_2$) increases allocative efficiency. Lowering the $\tilde{p}_2$ trigger price (i.e., decreasing the price collar between $\tilde{p}_1$ and $\tilde{p}_2$) has no significant effect on the level of allocative efficiency.*

Support. As shown in Table 2, allocative efficiency is very high in all treatments. Nevertheless, allocative efficiency is significantly higher ($p=0.05$) in LowP1 (97.4 percent on average) compared to Control (95.4 percent on average). However, there is no difference ($p=1.00$) in allocative efficiency between LowP2 (94.2 percent) and Control (95.39). The results reported in Column (4) of Table 3 are consistent with this.

We also find no impact of quantity changes on allocative efficiency.

Result 11. *Increasing or decreasing Q_1 and Q_2 has no significant impact on allocative efficiency.*

Support. Average allocative efficiency is between 95-96 percent in each of the Control, SmallQ2 and LargeQ2 treatments, and none of the differences are significant ($p>0.20$). Similarly, the coefficients on both LargeQ2 and SmallQ2 in Column (4) of Table 3 are insignificant.

4 Results: Individual Bidding Behavior

To conclude our analysis we now investigate individual bidding behavior. Within multi-unit auctions it is well known that bidders demand reduce (e.g., Ausubel et al., 2014): there is an incentive to bid at (or close to) their true value for the first allowance but bid below their value for the next allowance. The rationale is that if the buyers submit (near) truthful first bids they may stand a chance of obtaining an allowance but lowering the second bid means that they will pay a much lower price for all allowances obtained.

In our experiment, every subject places two bids within each round. These bids are constrained only in that they must be non-negative cent amounts.¹⁵ We categorize each person's bids as High and Low accordingly.¹⁶ We then compute the ratio of the bid to the value for each of these bids, with summary statistics in each treatment reported in Table 4. While clearly there is considerable heterogeneity in behavior, the bid-value ratio is slightly above one for the high bids, but approximately one for the low bids. Using Wilcoxon signed-rank tests, the high bid-value ratio is not significantly different from one in the LowP1 and LowP2 treatments, but is (above one) in the Control ($p=0.02$), LargeQ2 ($p=0.08$) and SmallQ2 ($p=0.06$). On the other hand, the low bid-value ratio is not significantly different from one in the Control, LowP2 and the LargeQ2 treatments, but is (below one) in the LowP1 ($p = 0.03$) and SmallQ2 treatments ($p = 0.03$). The differences between the high and low bid-to-value ratios are significant in all treatments ($p \leq 0.03$). These results provide indicative evidence that bidding is at least approximately consistent with what theory predicts.

¹⁵We exclude from our analysis in this section any bids that exceed 10; fortunately this comprises only 2.3% of our data.

¹⁶Overall, 72% of bids in the experiment involve different bids for the two units.

Table 4: Mean (Standard Deviation) of Bid-to-Value Ratios by Treatment

	Control	LowP2	LowP1	LargeQ2	SmallQ2	Overall
High Bid / Value	1.353 (5.859)	1.403 (3.216)	0.981 (0.622)	1.685 (8.297)	1.199 (2.228)	1.325 (4.903)
Low Bid / Value	1.056 (3.919)	1.139 (2.715)	0.850 (0.479)	1.384 (8.078)	0.882 (0.799)	1.066 (4.253)

We next consider the distribution of the bids in each treatment. Figure 6 shows the distribution of high bids, while Figure 7 shows the low bids. Also shown on both figures is the distribution of values, which, in accordance with the experimental design, is approximately uniform. Across all treatments, the bids disproportionately focus on the trigger prices, with the effects more pronounced with low values and in certain treatments such as SmallQ2. For example, more than 15% of low bids in SmallQ2 are close to $\tilde{p}_1$, triple the number of values that occur in that range. Thus subjects are routinely bidding at the trigger prices. Rationally, if a subject has a value close but below a trigger price it may be optimal to increase bids in order to increase the quantity of allowances sold while only marginally increasing the clearing price (thus potentially increasing their payoff). Responses from the survey questionnaire support the idea that subjects often use the trigger prices as focal points. Within written comments, 23% of subjects explicitly mentioned the trigger prices as a factor in their choice of bidding strategy.




Figure 6: Distribution of High Bids that Shows a Bimodal Distribution that Aligns with the Trigger Prices.




Figure 7: Distribution of Low Bids that Shows a Bimodal Distribution that Aligns with the Trigger prices.

It is important to note, however, that the bimodal distribution of bids and therefore clearing prices (Result 1) is only partially explained by bidder behavior. There is, in fact, a significant institutional effect on the determination of the clearing price. Figure 8 shows the predicted prices within each treatment if all bidders were to bid truthfully given their assigned private

values. Again, note that in such a case the distribution of predicted prices is bimodal and aligns with the two trigger prices. This suggests that the focal point nature of the trigger prices is explained by *both* bidder behavior and also the institutional design of the allowance reserves.




Figure 8: Predicted Clearing Price under Truthful Bids Generate a Biomodal Distribution that Aligns with the Trigger Prices.

It is likely that the larger the size of the allowance reserves relative to the initial cap may mean a greater density of clearing prices clustering around the trigger prices.¹⁷ Within our control treatment the size of both allowance reserves is one-half of the initial auction supply. While this is stylized, note that the allowance reserves within the RGGI design are also approximately one-half of the auction supply. Our evidence then suggests that the design of trigger prices and the size of the allowance reserves have important consequences for the outcomes of the regulatory system. Indeed the choice of trigger prices may have significant effects on the clearing price and the associated revenue and allocative efficiency.

5 Concluding Remarks

In the majority of pollution markets an increasing trend has been to allow for variable allowance supply. Adjusting allowance supply can assist firms with their abatement costs as well as spur innovation in pollution abatement technologies. A prominent design—used within the US Regional Greenhouse Gas Initiative (RGGI)—is the use of allowance reserves that can either increase or decrease the initial supply. Allowance reserves are fixed quantities of allowances that can either flow in or out of the regulatory system if the clearing price is sufficiently high or low: essentially a “soft” price floor and ceiling. These allowances are activated when the price reaches a “trigger price”, which is predetermined by the regulator. A proposed plan for 2021 will see RGGI create a dual allowance system, with both a lower and upper allowance reserve. Yet, while this novel regulatory design is forthcoming, there is sparse evidence on how dual allowance reserves will impact the regulatory outcomes.

The aim of this article is to provide an experimental analysis of dual allowance reserves, which underpins the proposed RGGI 2021 design. We find the choice of trigger prices has an important impact on the system: the trigger prices act as focal points, which is due to both institutional design and buyers’ bidding strategies. Thus the distribution of clearing prices is bimodal and aligns with the regulator’s two trigger prices. We find evidence that clearing

¹⁷Another exogenous parameter of influence is the variance of the value distribution. To operationalize buyers’ values we allowed the values to be randomly drawn from a finite distribution. An alternative would be, for example, to assume a normal distribution. While this may result in a larger number of realizations near the expectation and a lower amount at the extremes, a key parameter that will influence whether the allowance reserves were activated or not would be the variance of the distribution *relative* to the trigger price values (note that in our analysis, even without this assumption, clearing prices were very rarely below (above) the lower (upper) trigger price). Indeed if this distribution’s variance is relatively small then we may observe clearing prices that do not activate the allowance reserves. Conversely, a larger variance may result in a higher likelihood of observing clearing prices at the trigger prices.

prices increase in both trigger prices. We also find evidence that the trigger-price collar—the width between the two trigger prices—has important impacts on key regulatory criteria. In particular, a narrower trigger-price collar, increases revenue but decreases efficiency. We also test whether changes in quantities affect the regulatory system. While we show increasing the initial supply and decreasing the upper allowance reserve (holding total allowances constant) decreases the clearing price and increases allowances sold, we find that changes in quantities have no significant effect on revenue or efficiency. It follows that the regulatory design is more sensitive to the regulator’s chosen trigger prices than changes in the reserve quantities.

Our results provide concrete evidence that the regulatory outcome is sensitive to the regulator’s choice of trigger prices and allowance reserves. Depending on the regulator’s key objectives, we highlight the desirable design of trigger prices and allowance reserves for cost containment, revenue generation, and efficiency. The proposed RGGI scheme has fixed trigger price increases up to 2030 with the trigger-price collar widening year-on-year. Our evidence suggests, *ceteris paribus*, that the proposed scheme will place upward pressure on the clearing price, reduce the number of allowance sold, but increase the efficiency of the regulatory system. It may thus be more effective in containing emissions as compared to costs. Although our analysis is motivated by the new design features of RGGI, our analysis can provide general regulatory evidence for pollution markets embarking on variable supply policies. Future work could consider how alternative distributions of values impact the outcome of this regulatory system.

6 Appendix A: Experimental Instructions for the Control Treatment

Subject Instructions

Welcome to our experiment. Please read these instructions carefully. They are the same for every participant. Please do not talk with other participants and remain quiet during the entire experiment. Please turn off your mobile phone and don't switch it on until the end of the experiment. If you have any question, raise your hand and we will come to you.

The experiment will consist of 25 rounds. An auction with four bidders will take place in each round. You are one bidder, the other three are chosen randomly from the other people in the lab today. You will bid in the same group of four for the entire experiment today.

All money during this experiment is real money. At the end of the experiment, all of the money you have earned will be paid to you in Australian dollars. Your final payment will be rounded up to the nearest 20 cents.

Auction Rules

Each round consists of one auction. During each auction, the computer can sell **at most 6 units of the same good**. Each bidder can bid for **two** units of the good. The value of the good to each bidder will be different.

Each bidder's value for the good will be randomly drawn independently from the interval $0 \leq \text{value} \leq 10$; this means that every cent value between 0, 0.01, 0.02, ..., 9.98, 9.99, 10.00 has an equal chance of being selected. Both units of the good will have the same value for you.

Your value for the goods will not be revealed to the other bidders. Similarly, you will not know the values of the other bidders.

During each round, the auction proceeds according to the following rules: Each bidder makes an offer for each of the two units of the good. The offer reflects the maximum you are willing to pay for each unit of the good. The figure below shows a screenshot of the bidding screen.

Figure: Screenshot of auction decision screen

The screenshot displays a web-based bidding interface. At the top, it indicates 'Period: 1 of 2'. The main heading is 'Please enter your bid for each unit below'. Below this, there are three lines of small text: 'The computer can sell at most SIX units of the good.', 'The computer initially sells TWO units of the good.', and two conditional rules about the INTERIM price exceeding \$3.33 and \$6.66. The interface features a table with three columns: 'Unit number', 'Value', and 'Your bids'. For Unit 1, the value is 7.63, and for Unit 2, the value is also 7.63. Each unit has a corresponding blue input field for the bid. An 'OK' button is located at the bottom right of the screen.

Unit number	Value	Your bids
1	7.63	
2	7.63	

You will make a profit on each unit that you purchase for a price lower than your value.

Specifically, the *profit per unit* = *value of the unit* – *price paid*

For example, suppose your value for a unit is 4.80.

- If the price is \$4.00, then your profit per unit = $4.80 - 4.00 = \$0.80$.
 - If you purchase one unit then your profit is \$0.80.
 - If you purchase two units, then your profit is $0.80 \times 2 = \$1.60$
- If the price is \$6.00, then your profit per unit = $4.80 - 6.00 = -\$1.20$ (i.e. you make a loss of \$1.20 per unit).
 - If you purchase one unit then you make a loss of \$1.20.
 - If you purchase two units then you make a loss of \$2.40.

If you do not purchase an item then your profit for that round is zero.

At the end of the 25 auctions rounds, you will be paid your earnings in FIVE randomly selected rounds. These earnings will be added to your earnings in other parts of the experiment.

How is the Price Determined?

In each auction, the computer can sell up to 6 units of the good. The computer will initially supply 2 units of good, with the 4 additional units of supply held in reserve. These additional units will be supplied if the INTERIM price is high enough.

If the INTERIM price exceeds the first TRIGGER price of \$3.33, up to TWO additional units may be supplied. In this case, additional units of supply will be sold until i) the price falls to the TRIGGER price of \$3.33 or ii) the additional supply is exhausted.

The remaining TWO units will be supplied only if the INTERIM price exceeds the second TRIGGER price of \$6.66. In a similar way, additional units of supply will be sold until i) the price falls to the TRIGGER price of \$6.66 or ii) the additional supply is exhausted.

In this auction, all successful bidders pay the same price. Regardless of the final number of units supplied, the final price is always determined by the highest losing bid or the relevant trigger price, whichever is greatest.

The example on the following pages illustrates in detail how the price is determined.

After each participant has submitted bids for their two goods, the computer will determine the FINAL price as follows.

Step 1: Order the eight bids in each group from the highest to the lowest.

Refer to column (B) in the table below.

(A)	(B)
Bid Number	Ordered Bids
1	8.00
2	7.50
3	7.40
4	7.00
5	6.95
6	5.15
7	4.20
8	4.00

Step 2: Determine the INTERIM price for supplying all possible quantities.

The computer will sell between a minimum of 2 units of the good and a maximum of 6 units of the good. These possibilities are shown in column (C) of the table below.

For each possible quantity supplied, the INTERIM price is determined by the highest losing bid as shown in column (D) of the table below.

(A)	(B)	(C)	(D)
Bid Number	Ordered Bids	Possible Quantity Supplied	Interim Price
1	8.00		
2	7.50	2	7.40 ←
3	7.40	3	7.00
4	7.00	4	6.95
5	6.95	5	5.15 ←
6	5.15	6	4.20
7	4.20		
8	4.00		

If 2 units are supplied, the interim price equals the 3rd highest bid of \$7.40

If 5 units are supplied, the interim price equals the 6th highest bid of \$5.15

Step 3: Compare the INTERIM prices to the relevant TRIGGER price to determine the actual quantity sold

Remember the computer initially supplies 2 units of the good. If the INTERIM price exceeds the first TRIGGER price of \$3.33 then up to TWO additional units will be released. In column (E) we compare the INTERIM price with this first TRIGGER price of \$3.33. Notice that we make this comparison only for supply quantities of 2 and 3. (The other cells in the table are shaded to show they are not relevant.) This is because the computer only releases TWO additional units when the INTERIM price exceeds the first TRIGGER price of \$3.33.

(A)	(B)	(C)	(D)	(E)	(F)
Bid Number	Ordered Bids	Possible Quantity Supplied	Interim Price	Is Interim Price > \$3.33?	Is Interim Price > \$6.66?
1	8.00				
2	7.50	2	7.40	yes	
3	7.40	3	7.00	yes	
4	7.00	4	6.95		
5	6.95	5	5.15		
6	5.15	6	4.20		
7	4.20				
8	4.00				



Starting in the row where 2 units are supplied, we move down the rows until the first “no” entry in either column (E) or (F). In this example, when quantity equals 2, the interim price of \$7.40 > the first trigger price of \$3.33, so we enter “yes” in column (E), and continue to the next row. When the quantity equals 3, the interim price of \$7.00 > \$3.33, so we enter “yes” in column (E), and continue to the next row.

We then move to consider the row with a possible quantity supplied of 4. Now the relevant trigger price is \$6.66. See the table below. This is because the computer only releases the remaining TWO additional units if the INTERIM price exceeds the second trigger price of \$6.66. With a quantity of 4 supplied the interim price of \$6.95 > \$6.66 so we write “yes” in column (F) and move to the next row. However with 5 units supplied, the interim price is \$5.15 < \$6.66 so we write “no” and stop moving down the table.

(A)	(B)	(C)	(D)	(E)	(F)
Bid Number	Ordered Bids	Possible Quantity Supplied	Interim Price	Is Interim Price > \$3.33?	Is Interim Price > \$6.66?
1	8.00				
2	7.50	2	7.40	yes	
3	7.40	3	7.00	yes	
4	7.00	4	6.95		yes
5	6.95	5	5.15		no
6	5.15	6	4.20		
7	4.20				
8	4.00				



The total quantity sold in the auction is either:

- i) determined by the quantity supplied in the row where we first write “no”, or**
- ii) if we write yes in all rows, the quantity supplied equals six,**

In this example, the quantity supplied is 5 since we first wrote “no” in that row.

The final auction price is determined by either the interim price or the *relevant* trigger price, whichever is larger. *In this example, the interim price is \$5.15 and the relevant trigger price is \$6.66. The final auction price is the larger value of \$6.66*

Step 4: Allocate the goods

After the determining the quantity sold and the final auction price, the computer allocates the goods to the highest bidders. In this example, the five highest bids each purchase a good for the price of **\$6.66** as shown in columns (G) and (H) of the following table.

(A) Bid Number	(B) Ordered Bids	(C) Possible Quantity Supplied	(D) Interim Price	(E) Is Interim Price > \$3.33?	(F) Is Interim Price > \$6.66?	(G) Purchase	(H) Price Paid
1	8.00					Yes	6.66
2	7.50	2	7.40	yes		Yes	6.66
3	7.40	3	7.00	yes		Yes	6.66
4	7.00	4	6.95		yes	Yes	6.66
5	6.95	5	5.15		no	Yes	6.66
6	5.15	6	4.20			No	
7	4.20					No	
8	4.00					No	

Step 5: Determine the Payoffs of Successful Bidders

a) Suppose your two bids are 7.00 and 4.20 and your value is 7.88. Then you buy one unit and your payoff is $7.88 - 6.66 = 1.22$

b) Suppose your bids are 7.40 and 6.95 and your value is 5.66. Then you buy two units and your payoff is $(5.66 - 6.66) \times 2 = -2.00$ (i.e. a loss of 2.00).

Unsuccessful bidders earn a payoff of zero for that round.

Auction Results

At the end of each auction, participants will be informed of the eventual FINAL auction price and whether they have purchased any items. For each item you purchase, you will earn the difference between the value of the good minus what you paid for the good, so you will earn profits on each good you purchase at a FINAL price less than your value.

Note that only you will know *your* actual bids, allocations, and payoffs in each round. A sample feedback screen is shown below.

Figure: Screenshot of the Feedback screen

The screenshot displays a feedback screen titled "Auction Outcome". It is divided into two main sections. The top section shows aggregate results: "Total number of units sold" is 4 and "Auction price" is 4.00. The bottom section shows individual results: "Number of units you purchased" is 1, "Your value" is 7.30, and "Your earnings this period" is 3.30. A red button labeled "Next Period" is located in the bottom right corner.

Auction Outcome	
Total number of units sold	4
Auction price	4.00

Number of units you purchased	1
Your value	7.30
Your earnings this period	3.30

Next Period

Auction Sets

After this, a new round starts. In the new auction, you will continue to bid in the same group of four bidders as in the previous round. At the beginning of each round, a new value will be randomly selected for each participant as described earlier. Overall, there are 25 auctions.

At the end of the 25th round, please answer a short questionnaire that appears on your computer screen.

After that, we will randomly select FIVE rounds and you will be paid your auction earnings from those five rounds.

Finally, note that the rules are the same for every participant. If you have any questions, raise your hand and we will come to you. Before starting the experiment, we will ask you some questions to check that you understand these instructions. You will earn \$0.50 for each correctly answered question.

Note that any ties will be broken randomly.

Summary

- You will be bidding in 25 auctions in fixed groups of four bidders.
- You will be bidding for two units of the good with your value for the goods randomly drawn from a uniform distribution between 0 and 10.
- This value will be randomly redrawn each round. The value chosen is completely independent of the value chosen in any previous rounds.
- The computer will initially supply two units of the good with four additional units held in reserve.
- Two additional units will be released only if the interim price exceeds the first trigger price of \$3.33.
- The final two units will be released only if the interim price exceeds the second trigger price of \$6.66.
- The interim price is determined by the highest losing bid.
- The goods available will be sold to the highest bidders who all pay the same price.
- You will earn profit (or loss) = your value – price paid for each good you purchase.

References

- Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M. and Weretka, M. (2014). 'Demand reduction and inefficiency in multi-unit auctions', *The Review of Economic Studies*, **81**, pp. 1366–1400.
- Back, K. and Zender, J. F. (2001). 'Auctions of divisible goods with endogenous supply', *Economics Letters*, **73**, pp. 29–34.
- Burtraw, D., Palmer, K. and Kahn, D. (2010). 'A symmetric safety valve', *Energy Policy*, **38**, pp. 4921 – 4932.
- Burtraw, D., Charlie Holt and, K. P., Paul, A. and Shobe, W. (2018). 'Quantities with prices', In *Resources for the Future Discussion Paper 18-08*.
- Coursey, D. L. and Smith, V. L. (1983). 'Price controls in a posted offer market', *The American Economic Review*, **73**, pp. 218–221.
- Damianov, D. S. and Becker, J. G. (2010). 'Auctions with variable supply: Uniform price versus discriminatory', *European Economic Review*, **54**, pp. 571–593.
- Eckel, C. C. and Grossman, P. J. (2008). 'Forecasting risk attitudes: An experimental study using actual and forecast gamble choices', *Journal of Economic Behavior & Organization*, **68**, pp. 1 – 17.
- Engelbrecht-Wiggans, R., List, J. A. and Reiley, D. H. (2006). 'Demand reduction in multi-unit auctions with varying numbers of bidders: Theory and evidence from a field experiment*', *International Economic Review*, **47**, pp. 203–231.
- Engelmann, D. and Grimm, V. (2009). 'Bidding Behaviour in Multi-Unit Auctions—an Experimental Investigation', *The Economic Journal*, **119**, pp. 855–882.
- Engelmann, D. and Müller, W. (2011). 'Collusion through price ceilings? In search of a focal-point effect', *Journal of Economic Behavior & Organization*, **79**, pp. 291 – 302.
- Fell, H., Burtraw, D., Morgenstern, R. D. and Palmer, K. L. (2012). 'Soft and hard price collars in a cap-and-trade system: A comparative analysis', *Journal of Environmental Economics and Management*, **64**, pp. 183 – 198.
- Fell, H. and Morgenstern, R. (2010). 'Alternative approaches to cost containment in a cap-and-trade system', *Environmental and Resource Economics*, **47**, pp. 275–297.
- Fischbacher, U. (2007). 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, **10**, pp. 171–178.

- Gode, D. D. K. and Sunder, S. (2004). 'Double auction dynamics: structural effects of non-binding price controls', *Journal of Economic Dynamics and Control*, **28**, pp. 1707 – 1731.
- Greiner, B. (2015). 'Subject pool recruitment procedures: organizing experiments with orsee', *Journal of the Economic Science Association*, **1**, pp. 114–125.
- Grüll, G. and Taschini, L. (2011). 'Cap-and-trade properties under different hybrid scheme designs', *Journal of Environmental Economics and Management*, **61**, pp. 107 – 118.
- Holt, C. A. and Shobe, W. M. (2016). 'Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the EU ETS market stability reserve', *Journal of Environmental Economics and Management*, **80**, pp. 69 – 86.
- Isaac, R. M. and Plott, C. R. (1981). 'Price controls and the behavior of auction markets: An experimental examination', *The American Economic Review*, **71**, pp. 448–459.
- Khezr, P. and MacKenzie, I. A. (2018). 'Permit market auctions with allowance reserves', *International Journal of Industrial Organization*, **61**, pp. 283 – 306.
- Knittel, C. R. and Stango, V. (2003). 'Price ceilings as focal points for tacit collusion: Evidence from credit cards', *American Economic Review*, **93**, pp. 1703–1729.
- Kollenberg, S. and Taschini, L. (2016). 'Emissions trading systems with cap adjustments', *Journal of Environmental Economics and Management*, **80**, pp. 20 – 36.
- LiCalzi, M. and Pavan, A. (2005). 'Tilting the supply schedule to enhance competition in uniform-price auctions', *European Economic Review*, **49**, pp. 227–250.
- Lopomo, G., Marx, L. M., McAdams, D. and Murray, B. (2011). 'Carbon allowance auction design: An assessment of options for the United States', *Review of Environmental Economics and Policy*, **5**, pp. 25–43.
- McAdams, D. (2007). 'Adjustable supply in uniform price auctions: Non-commitment as a strategic tool', *Economics Letters*, **95**, pp. 48–53.
- Mckinsey and Company (2016). 'Greenhouse gas abatement cost curves', , Mckinsey and Company. <http://www.mckinsey.com/business-functions/sustainability-and-resource-productivity/our-insights/greenhouse-gas-abatement-cost-curves>.
- Murray, B. C., Newell, R. G. and Pizer, W. A. (2009). 'Balancing cost and emissions certainty: An allowance reserve for cap-and-trade', *Review of Environmental Economics and Policy*, **3**, pp. 84–103.

- Pagnozzi, M. and Saral, K. J. (2017). 'Demand Reduction in Multi-Object Auctions with Resale: An Experimental Analysis', *The Economic Journal*, **127**, pp. 2702–2729.
- Perino, G. and Willner, M. (2016). 'Procrastinating reform: The impact of the market stability reserve on the EU ETS', *Journal of Environmental Economics and Management*, **80**, pp. 37 – 52.
- Perkis, D. F., Cason, T. N. and Tyner, W. E. (2016). 'An experimental investigation of hard and soft price ceilings in emissions permit markets', *Environmental and Resource Economics*, **63**, pp. 703–718.
- Rey, P. and Tirole, J. (2019). 'Price caps as welfare-enhancing coopetition', *Journal of Political Economy*, **127**, pp. 3018–3069.
- RGGI (2013). 'Summary of RGGI model rule changes'. available at http://rggi.org/docs/ProgramReview/_FinalProgramReviewMaterials/Model_Rule_Summary.pdf.
- RGGI (2017). 'RGGI 2016 Program Review: Principles to accompany model rule amendments'. available at https://www.rggi.org/sites/default/files/Uploads/Program-Review/12-19-2017/Principles_Accompanying_Model_Rule.pdf.
- Roberts, M. J. and Spence, M. (1976). 'Effluent charges and licenses under uncertainty', *Journal of Public Economics*, **5**, pp. 193 – 208.
- Salant, S., Shobe, W. and Uler, N. (2019). 'The effects of non-binding price floors on the price of storable assets'. University of Maryland, Mimeo.
- Scherer, F. and Ross, D. (1990). *Industrial Market Structure and Economic Performance*, third ed., Boston.: Houghton Mifflin, Boston, 3rd edition.
- Shobe, W., Palmer, K., Myers, E., Holt, C., Goeree, J. and Burtraw, D. (2010). 'An experimental analysis of auctioning emission allowances under a loose cap', *Agricultural and Resource Economics Review*, **39**, pp. 162–175.
- Shobe, W., Holt, C. and Huetteman, T. (2014). 'Elements of emission market design: An experimental analysis of California's market for greenhouse gas allowances', *Journal of Economic Behavior & Organization*, pp. 402 – 420.
- Stranlund, J. K., Murphy, J. J. and Spraggon, J. M. (2011). 'An experimental analysis of compliance in dynamic emissions markets', *Journal of Environmental Economics and Management*, **62**, pp. 414 – 429.

Stranlund, J. K., Murphy, J. J. and Spraggon, J. M. (2014). 'Price controls and banking in emissions trading: An experimental evaluation', *Journal of Environmental Economics and Management*, **68**, pp. 71 – 86.