July 15, 2022

Dear Sir or Madam,

We are pleased to share the accompanying comments to the Washington State Department of Ecology on the proposed rule for the Climate Commitment Act Program, WAC Chapter 173-446.

RFF is an independent, nonprofit research institution in Washington, DC. Its mission is to improve environmental, energy, and natural resource decisions through impartial economic research and policy engagement. RFF is committed to being the most widely trusted source of research insights and policy solutions leading to a healthy environment and a thriving economy.

While RFF researchers are encouraged to offer their expertise to inform policy decisions, the views expressed here are those of the individual authors and may differ from those of other RFF experts, its officers, or its directors. RFF does not take positions on specific policy proposals.

The views expressed by Dr. Shobe are his own and do not necessarily represent the position of the Rector and Board of Visitors of the University of Virginia.

If you have any questions or would like additional information, please contact us at the email addresses below. Any references cited are available from the authors.

Sincerely,

Dallas Burtraw (burtraw@rff.org)
William Shobe (shobe@virginia.edu)

Comments on the proposed WAC 173-446: Climate Commitment Act Program Rule

Dallas Burtraw, Resources for the Future
William Shobe, University of Virginia
July 15, 2022

We appreciate the opportunity to provide comments on the proposed rule on the Climate Commitment Act Program, WAC Chapter 173-446. Our comments will focus primarily on the implementation of the emission containment reserve (ECR) as discussed in sections 220, 300, 340, 357, 370 and 375.

In its standard implementation, an ECR is a mechanism for automatically adjusting the supply of emission allowances under conditions where the price of allowances is below that anticipated at the program's outset. Hence, the ECR acts to accelerate emissions reductions when the market price signals that it is inexpensive to do so.

A properly designed ECR offers a rule-based approach to adjusting allowance supply in response to market signals about allowance scarcity. This reduces uncertainty for both compliance entities and lowers administrative costs for regulators. In contrast, administrative adjustments to supply can propagate regulatory uncertainty and the expectation that one administrative intervention may foreshadow other additional program interventions.

Establishing a functional ECR at the initial implementation of the program sets market expectations for the long run and helps ensure the durability of the program. Long experience in many regulatory settings demonstrates that from an administrative perspective it is easier to establish program features at the outset of a program that anticipate potential future concerns than to adjust program design in response to concerns that may arise in the moment. Implementing the ECR in Washington before the rest of the WCI does not harm Washington's interests. It does have the potential to push the WCI towards a better market design with Washington as first mover.

Further, the ECR design in the current rule can be improved to maximize its benefits. A simpler approach than exists in the current regulation would implement the ECR as a reserve price in the primary auction, and any reserved allowances could be retired or, alternatively, placed in equal shares in the APCR tiers. This design would strengthen the ability of the ECR to reduce excess price volatility (and market uncertainty), improve price discovery, and simplify program administration.

We make three main points about the ECR proposal in the proposed rule:

A. We strongly encourage the agency to include ECR provisions in the rule and set an ECR trigger price. This should occur whether other revisions to the proposed regulation are adopted or not.

B. The proposed design of the ECR treats the ECR as a separate account holding a stock of allowances for re-allocation back into the market by grant or by supplemental auction. The proposed design could be greatly enhanced if it were implemented in a hybridized way by designating a portion of allowances in the ECR for distribution to allocation to energy intensive trade exposed (EITE) entities and new entrants as required by statute, and distributing the remainder of the allowances within the quarterly auction framework.
C. Implementing the ECR through a reserve (trigger) price in the quarterly auction would simplify administration and embody best practice and deliver the maximum benefits.

The remainder of these comments provide rationale for these suggestions.

A) Washington should include an ECR and set an ECR trigger price at the beginning of the program.

An important evolution in the design of emissions markets is the move away from specifically fixed emissions allowance caps to allowance supply schedules that respond to the equilibrium price identified in an allowance auction. This reform helps to remedy the interaction of carbon pricing with other regulatory programs while retaining the virtue of price discovery and cost effectiveness associated with carbon markets.

A concern of many stakeholders is uncertainty about allowance prices, and that prices may be higher than anticipated, which is understandable given unfamiliarity with the program. The cost containment elements of the program are designed to ameliorate this concern. Very high prices can be prevented by making some additional number of allowances available at price trigger points. In the proposed rule, this is accomplished by the Allowance Price Containment Reserve (APCR) and the price containment units, both of which increase the allowance supply if the market is tighter than expected. In sum, the implementation of the ECR has little relevance to stakeholder concerns initially, because those concerns are about very high prices.

However, the ECR has important value in shaping price expectations for the long run. In every important market for atmosphere resources (sulfur dioxide, nitrogen oxides and carbon dioxide) in North America and Europe, after initial price volatility representing uncertainty and hedging activities, prices have fallen to below expectations and often fallen in real terms.[1] Perhaps surprisingly, these periods of low prices rather than high prices have constituted the major challenge to the durability of these programs. In the long run, the interaction of the carbon market with other regulatory programs becomes important as compliance entities make investments that anticipate the state's long-term climate goals, and which are informed by the current and anticipated future carbon price. The ECR provides a guardrail against unexpected price declines, including potential price effects that may result from interactions of the carbon market with companion regulatory policies.

[1] "Recognizing Gravity as a Strong Force in Atmosphere Emissions Markets," 2018 (Dallas Burtraw and Amelia Keyes), *Agricultural and Resource Economics Review,* 47(2): 201-219.

To guard against extreme price declines, it has become usual practice to have a reserve price in the allowance auction, which provides a price floor in the auction.[2] This means that the total number of allowances available responds to the market demand for allowances, just as one observes in commodity markets. The ECR adds a second reserve price, set at a level above the auction price floor, that applies to 10 percent of the allowances available for sale. Importantly, the ECR lowers price volatility by making automatic adjustments to the long-run allowance supply. This adjustment helps stabilize auction proceeds for program-related investments.

The performance of an ECR has been shown theoretically to improve emission market performance, and these results have been borne out in simulations, experiments, and actual practice.[3] The ECR is a design element of the Regional Greenhouse Gas Initiative (RGGI) program beginning in 2021 and has been identified by observers and researchers as a meaningful reform elsewhere.[4]

We believe the time to implement an ECR is when it is not expected to be immediately relevant, which based on experience in other programs is likely to be at the outset of the program. Hence, we believe the ECR trigger price should be set in Section 340 rather than suspended as in the proposed rule. The implementation of the ECR trigger price does not generate any disadvantages for Washington, but rather protects Washington's emission market against unexpectedly low emission prices such as has occurred during later stages in a number of previous emission markets.

Washington's ECR would set an important precedent for other states potentially joining in a regional emissions market. If at some point, suspension of the ECR trigger price is required to enable program linkage, then action could be taken at that time. In the meantime, the presence of an ECR sets expectations for discussions across jurisdictions and provides a positive example that could propagate to other jurisdictions and strengthen climate policy generally.

Summary: Given the strong evidence in favor of using an ECR, we believe that it is very important that this feature be included in the proposed rule. WAC 173-446-340 should be changed so that it instates a trigger price from the outset.

---

[2] The EU Emissions Trading System has implemented a different mechanism called the Market Stability Reserve to accomplish similar goals.

[3] See: "Price-Responsive Allowance Supply in Emissions Markets," 2022 (Dallas Burtraw, Charles Holt, Karen Palmer, and William Shobe). *Journal of the Association of Environmental and Resource Economics, 9 (5): 851–884, https://doi.org/10.7910/DVN/DHU5PM)* , and Roberts, M.J., and M. Spence. 1976. "Effluent Charges and Licenses Under Uncertainty." *Journal of Public Economics*, 5 (3-4): 193-208.

[4] 2021 IEMAC Annual Report | CalEPA

B) Adjustments to the ECR design can greatly improve its performance by integrating its operation within the quarterly auction framework, and separating out other functions including allocation to EITE entities and new entrants

The purpose of the ECR is to accelerate emissions reductions when it is inexpensive to do so, and to reduce unnecessary market uncertainty (price volatility). Maintaining price stability will enhance the availability of auction proceeds directed at investments under the program. But in markets for commodities like emission allowances, the price is a reflection of expectations about the long-run balance between supply and demand. Shifting the availability between periods or among market participants will not have significant effects on those expectations of scarcity, and hence will not have a significant or any effect on price volatility. To be most effective at reducing market uncertainty, the ECR must be designed to adjust long-run supply.

Unfortunately, the ECR implementation in the proposed rule does little or nothing to address long-run imbalances between supply and demand. Even if a low market price clearly signals low-cost emissions reduction opportunities and an excess supply of allowances, the proposed ECR does not appreciably change the number of allowances allocated. This is because many, if not all, of the allowances sequestered in the ECR account are promptly recycled back into the allowance supply through supplemental auctions. Because the long-run supply doesn't change, any effect on current prices will be smaller than what is needed, indeed if there is any effect at all.

We offer a detailed description of the proposed regulation and examples of potential outcomes. Section 375 of the proposed rule specifies two avenues for the distribution of allowances held in the ECR: (1) free distribution to EITE facilities and (2) an auction to covered entities and opt-in entities whenever a new covered or opt-in entity enters the program. Figure 1 depicts the various flows into and out of the ECR. The first problem with this language is that there is no ordering or priority given to these two purposes, which may conflict. The second problem, and the one most relevant for the effectiveness of the ECR is that these allowances are simply recycled from the primary auction to a secondary auction. This mechanism does not reliably reduce the excess supply of emissions. These provisions should have no effect on the market price.

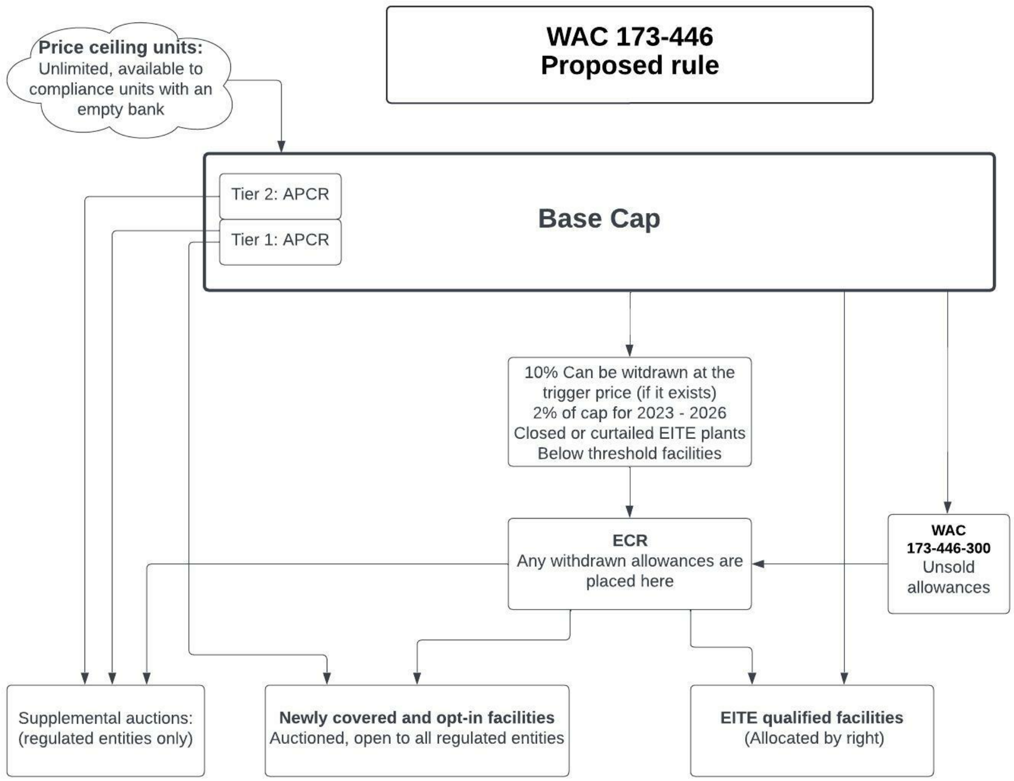

Figure 1: Allowance allocation in the proposed rule

A specific example may help clarify this issue. Suppose a quarterly auction closes at or below a $25 ECR trigger price, resulting in 10 allowances not being sold and being placed into the ECR. Suppose that a EITE distribution of 20 allowances is required. That distribution will draw down the ECR by 10 and subsequently it will draw down the next auction amount by 10. Had the ECR not been triggered withholding the sale of the 10 allowances in the first auction, the next auction would be reduced by 20. In either case, there has been a net transfer of 20 allowances from the available stock to the EITE facility. The scarcity of allowances has not changed. No effect on price should be observed due to the presence or absence of the ECR.

Alternatively, suppose that there is no call for distributions to EITE facilities, but a new or opt-in facility triggers the auctioning of the ECR stock of 10 allowances. The same compliance entities that bid in the quarterly auction will bid in this auction. The 10 allowances removed earlier in the primary auction would be brought into the market. Although the regulatory language appears ambiguous, it appears that because all auctions use the same structure including an ECR trigger price applied to ten percent of the allowances for sale in that auction.[5] Consequently, in the secondary auction, if market fundamentals have not changed, 90 percent

[5] See WAC 173-446-357. Note that, even if the trigger reserve applied to all ECR allowances sold in the supplemental auctions, the proposed mechanism would not achieve the intended result because the ECR mechanism would not adjust the long-run supply of allowances.

of the ECR stock will re-enter the market and only 10 percent of the ECR stock will remain in the ECR. The ECR allowances will be sold at a market price below the trigger price. In the worst case, anticipation of this possible outcome could influence the behavior of market participants. The potential for mischief is great. One immediate (if incomplete) fix to this possible outcome is to apply the ECR trigger price to all allowances sold in the supplemental auction.

Alternatively, if the new source is large enough to drive the market price above the ECR trigger price, the new auction closing price will rise above the trigger. Either way, the presence or absence of the ECR has had little or no effect on the stock of allowances in the market. Thus, we cannot expect the market price to be influenced in a meaningful way due to the presence of the ECR.

The situation in which a new source triggers a withdrawal from the ECR may be a somewhat rare occurrence, but the ECR stock must still be considered to be part of the long-run supply of allowances, and its presence will put downward pressure on the market price. The entire ECR stock (or at least 90% of it) could reenter the market in a single auction triggered by a single new entrant, potentially at a price below the ECR trigger price.

The sequence of auctions provided for in the current draft rule may create an opportunity to benefit from strategic behavior, affecting the likelihood these situations are observed. The fundamental problem with the current ECR proposal is that it is implemented as a temporary separate account where allowances sit for a short time before reentering the market, and the allowances directed to the benefit of EITE and new facilities are comingled with other allowances in the ECR. Consequently, it is possible that allowances flow into the ECR if demand is slack but then flow right back out into the same slack market. The anticipation of the future return of ECR allowances to the market must lower the current market price.

The proliferation of auctions with different allowed participation may contribute to unnecessary price volatility. Various auctions in the proposed rule with different numbers of participants and different rules about how many allowances parties may purchase can be expected to result in different prices in different auctions, differences that have little to do with underlying market expectations about allowance scarcity. It may also create incentives for market participants to manipulate their bids in ways that are very hard to predict in advance. The better option is to reduce the number and variety of auctions.

Summary: The current design of the ECR treats the ECR as a separate account holding a stock of allowances for re-allocation back into the market by grant or by supplemental auction. If withdrawals are triggered to benefit EITE or new facilities, the design has very little chance of reducing price volatility.

C) A simple adjustment to the proposed rule can distinguish the distribution of allowances to benefit EITE and new entrants from other (general) allowances in the ECR. The ECR allowances not used for EITE facilities and new entrants would be sold in the normal quarterly auction but with the trigger price as the reserve price for those allowances, as currently provided in Section 357. This would embody best practice and deliver maximum benefits

The function of the ECR in the proposed rule can be made consistent with best practice design with a small change in the proposal. Instead of defining the ECR as a separate account into which allowances are placed for later sale at auction, the ECR can be defined as the 10 percent of allowances that can be removed from any allowance auction at the ECR trigger price. (This is equivalent to applying an ECR reserve price to the 10 percent of allowances.) Before the auction, the number of allowances subject to this trigger reserve price would be reduced by any *required distribution to EITE facilities*. If legislation requires new and opt-in facilities to have preferential access to the ECR, they could be offered ECR allowances at the trigger price, which guarantees they will receive the allowances. We should emphasize that the best approach both for climate and for supporting good market function would be to separate ECR auction design from the EITE distribution and the distribution to new sources.

Any auctioned ECR allowances not meeting the (trigger) reserve price would be retired. Another, somewhat less preferred, option would be to add unsold allowances to the APCR. Retirement is preferred because it better advances climate action and better fulfills the intended purpose of an ECR by adjusting the allowance supply. Retiring allowances that do not meet a reserve price is equivalent to adding them to the stock of price containment units, what we might think of as Tier 3 of the ACPR. Allowances not sold at the regular auction reserve price would be reallocated in the same way. Figure 2 shows the allowance flows in this suggested approach to the ECR.

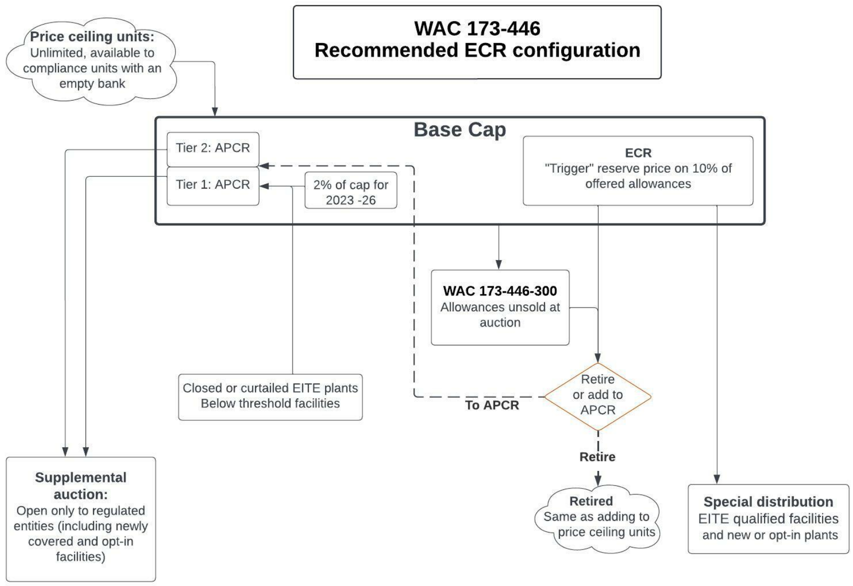

Figure 2: Recommended ECR configuration

This mechanism responds to the slack allowance market by taking the ECR allowances and making them available only if the market becomes tight enough that the price rises high enough to make it worth purchasing them at the higher prices in the APCR or price ceiling. The key feature of the ECR design is to add steps to the allowance supply so that, as the market becomes over-supplied, the supply automatically shrinks. The portion of supply removed is then made available in the eventuality of high future scarcity. If companion regulatory policies are effective enough in reducing emissions so that the price never reaches the price at which more allowances are released, it is conceivable that these allowances would never be needed, but they serve as valuable insurance against potential shocks resulting in unexpectedly high prices

This is how the ECR is implemented in RGGI; the RGGI ECR is implemented as a reserve price in the primary auction and if the auction clearing price is at or below the ECR trigger (reserve) price then some portion or all of the ECR allowances are not sold. This design has the considerable advantage of reducing the complexity of the auction provisions in this proposal. The only auctions needed in this revised ECR are the regular periodic auctions and any sales of allowances from the APCR and price containment units.

EITE facilities receive their allowances by right from the total available amount to be auctioned (or possibly from the ECR portion of the allowances at auction), so these allowances would not be available at auction. New, expanded and opt-in facilities would also receive their allocation from the total auction quantity (or, if necessary, from the ECR portion), offered at the trigger price. Administration would be simplified and competition improved by reducing the number of supplementary auctions.

Further, we suggest the proposed rule should minimize the instances where participation in an auction or market activity is limited to compliance entities. Such rules invite costly activities on the part of brokers, investors, and compliance entities to circumvent them and are very difficult to enforce. Moreover, emission markets are generally quite liquid, hence the market price of allowances and the price of allowances at auction will be quite close, so there is little to no effect resulting from restricting participation in the auction even if it can be effectively enforced. However, doing so may convey a disadvantage to smaller compliance entities with less in-house market expertise because brokers often provide valuable services to compliance entities. The participation of brokers and investors in auctions generally should be encouraged.

Summary: A small modification to the proposed rule can greatly strengthen the ECR provisions. In Section 375, remove the reference to the ECR being a separate account. The rule can simply provide that allowances that are not sold be retired. No supplemental auctions relating to implementation of the ECR, such as those in Paragraph 2, need be mentioned. This small change will greatly enhance the function of the ECR provisions in this proposed rule.

Price-Responsive Allowance Supply in Emissions Markets

Dallas Burtraw, Charles Holt, Karen Palmer, William Shobe

Abstract: Environmental policy with uncertainty is often posed as a choice between price and quantity instruments. Adding flexibility to fixed policy instruments can improve outcomes. Roberts and Spence noted the efficiency advantages of matching emissions allowances supply to the marginal damage schedule. We propose an implementable approach to making that match, an approach we call "price-responsive supply," which treats prices and quantities as simultaneously determined in the allowance auction. For competitive environments, price-responsive supply outperforms fixed-price and fixed-quantity instruments. Price-responsive supply can enhance the performance of real-world regulatory environments through an automatic adjustment mechanism that responds instantaneously to new information about abatement costs. We demonstrate the improved performance of price-responsive supply in experiments and simulations. A price-responsive supply schedule, while offering efficiency advantages, also translates the cost-lowering effects of other, coincident policies into accelerated reductions under an emissions cap, thereby helping to resolve the waterbed effect.

JEL Codes: Q48, Q54, Q58

Keywords: cap and trade, climate policy, greenhouse gas, climate change, electricity

THE QUESTION OF WHETHER TO USE PRICES (emissions charges) or quantities (tradable emissions allowances) to implement incentive-based regulation is one of the most persistent debates in the environmental economics literature (Montgomery 1972;

Dallas Burtraw (burtraw@rff.org) is at Resources for the Future, Washington, DC. Charles Holt (cah2k@virginia.edu) is at the University of Virginia. Karen Palmer (palmer@rff.org) is at Resources for the Future, Washington, DC. William Shobe (shobe@virginia.edu) is at the University of Virginia. We appreciate comments from Lori Bennear, Marty Weitzman, participants in the University of Virginia Experimental Economics Reading Group, Sophia Lochner, and three helpful anonymous reviewers. Alex Watkins made significant contributions to sec. 2 of this paper, Anthony Paul and Nicholas Roy made significant contributions to sec. 4. We are grateful to Georgia Beazley

Dataverse data: https://doi.org/10.7910/DVN/DHU5PM

Received August 31, 2020; Accepted April 29, 2022; Published online June 27, 2022.

Journal of the Association of Environmental and Resource Economists, volume 9, number 5, September 2022.
 Published by The University of Chicago Press for The Association of Environmental and Resource Economists.
https://doi.org/10.1086/720690

Weitzman 1974). A related literature has discussed "hybrid" systems, where features of price and quantity instruments are combined using quantity triggers or price collars. The question is motivated by the search for an efficient instrument in the presence of quantifiable uncertainty about benefits and costs. We propose to replace the "prices *versus* quantities" perspective with one that unifies the two approaches into a price-responsive allowance supply—"prices *with* quantities"—a feasible instrument that simultaneously determines price and quantity in an auction, implementing an upward-sloping emissions allowance supply curve rather than a fixed price or fixed quantity.

Building on the Roberts and Spence (1976) description of a hybrid system, this approach starts with the recognition that prices and quantities are jointly determined outcomes of an economic process. They showed that, under uncertainty, making emissions allowances available along the current best estimate of the marginal damage schedule could achieve an efficient expected outcome. Implementing this approach has not been considered practical previously,[1] but the increasing use of auctions for allocating emissions allowances enables a dynamic adjustment of allowance quantity using reserve prices to match the expected marginal damage curve, as originally suggested by Roberts and Spence. Price floors and caps using reserve prices add guardrails that can be represented as horizontal segments of an otherwise vertical auction supply function. Price-responsive supply adds an arbitrary number of steps or even a positively sloped segment to the supply function that is used to settle allowance auctions.

Considering this new view of the emissions allowance supply as automatically adjusting to new information about abatement costs, we argue that emissions markets are inherently a more flexible instrument than a feasible emissions charge. We demonstrate experimentally how price-responsive supply can be implemented and that it can reduce price and revenue volatility. Using the examples of the Regional Greenhouse Gas Initiative (RGGI) and the Western Climate Initiative, we argue that initial steps toward implementing elements of price-responsive supply have been important contributors to the successful operation of these cap-and-trade regimes.

Along with the technical aspects of instrument choice, policy makers face several obstacles to implementing an economic policy instrument to regulate greenhouse gas emissions that contribute to climate change. First among these obstacles is that calibrating carbon prices at levels high enough to approximate marginal climate damages has not been generally politically achievable. One among many persistent reasons is the challenge of

---

and Maya Domeshek for research assistance. The research was supported by the Energy Foundation, the Merck Family Fund, Mistra Carbon Exit, the National Science Foundation (SBR 1459918), and Resources for the Future's Research and Policy Engagement Program.

1. Roberts and Spence noted the problem of increasingly thin markets in their proposed implementation. Baumol and Oates (1988) dismissed the idea of price floors because previous programs had used free allocations, which suggests a contingent property right that would have to be taken away or purchased by the government if prices fell and allowances were retired.

economic competition from firms in unregulated jurisdictions. Disparate emissions pricing policies across jurisdictions cause economic and emissions leakage, constraining politically sustainable price levels.

We start with the recognition that policy makers may be prevented from approaching the problem purely as a matter of choosing the technical economic optimum that balances marginal climate damages with marginal abatement costs. One can imagine an adjusted marginal cost function that includes political and economic costs, which bound the policy maker away from the usual representation of the economic optimum. In this more complete representation of the policy maker's problem, one might expect policy to be at least partially oriented toward influencing the factors that create a difference between technical marginal costs and adjusted marginal costs as they are perceived in the policy process.[2] In fact, the adjusted marginal cost function is responsive to investments in technological change and infrastructure to lower marginal costs and policy coordination among jurisdictions to reduce leakage.[3]

The pricing of greenhouse gas emissions, in particular, faces a special challenge due to the subsidiary structure of governance under which local governments and civic organizations often take measures to reduce emissions. Under a quantity-based emissions cap, these measures do not yield additional emissions reductions and instead reduce the market price of the quantity instrument (emissions allowance) which reduces the price-based incentive for regulated parties to achieve emissions reductions. This phenomenon, known as the waterbed effect, would not be observed under an emissions tax.[4] But, in the political realm, advocates of climate policy may not see tangible assurance of emissions reductions under a carbon tax, as one might under an emissions cap that is declining over time. For this and other reasons,[5] emissions markets are more widely employed than taxes in climate policy. Meanwhile subsidiary actions intended to achieve reductions that are additional to those achieved by the market are common: the *Long-Term Strategy of the United States* (2021) describes four strategic pillars to achieving US climate goals, three of which (innovation, nonfederal leadership, and all-of-society action) are additional to federal leadership.

2. "We should not shoehorn the (climate policy) problem into familiar (economic) structures just because they are familiar. That approach simply fails to capture the issues at stake" (Oswald and Stern 2019).

3. There may even be a general opposition to the very idea of having politicians or a government agency setting prices.

4. The waterbed effect refers to the rebound from efforts to reduce emissions at one location that results in an increase in emissions at another location when total emissions are capped (Perino 2018).

5. For example, in many jurisdictions, emissions trading can be implemented under existing regulatory authority, but carbon taxes require legislative approval. Emissions trading programs have introduced measures to limit leakage and address competitiveness that do not yet have practical analogues implemented under emissions taxes.

When, via the waterbed effect, the market design effectively preempts the influence of additional efforts by subsidiary actors, it undermines strategic planning and political support for carbon pricing (Fankhauser et al. 2010; Dolan 2021).

Furthermore, in most cases, and especially in the context of climate policy, the regulator's problem is not a one-time choice of instrument under uncertainty but instead a sequence of policy choices selecting from an evolving set of achievable policy options. In practice, policy makers seeking to capture the promise of cost-effectiveness from carbon pricing have pursued an incremental approach. In a dynamic policy context, the climate policies we observe tend to follow a pattern foreshadowed in Baumol and Oates (1971), where regulators set a science-based target and chart a contingent pathway to that target. Everywhere, we observe that policy makers invoke a variety of "companion" policies to promote new technology and cultivate political alliances, aiming indirectly at reducing greenhouse gas emissions. Frequently, these policies are supported by investments of auction proceeds from the trading program. Over time, these activities contribute to changes in economics, politics, and technology that collectively influence the set of feasible policy options.[6]

This paper offers a novel contribution by reframing the conventional instrument choice for achieving an economic optimum from one of price versus quantity to one of implementing a single instrument that reduces the costs of uncertainty by reducing the frictions inherent in incremental policy responses to new information. Price-responsive allowance supply (as opposed to emissions prices or quantities alone) can help resolve the waterbed dilemma and enable the dynamic incremental process described by Baumol and Oates that can move policy toward a long-term target cost-effectively. Unlike quantity adjustment mechanisms for emissions taxes that adjust ex post in response to observed emissions, and price collars in trading programs that offer only guardrails on prices, a price-responsive allowance supply provides immediate and continuously responsive feedback to decision makers.

We investigate the performance of price-responsive supply functions in an experimental setting and in simulation modeling. We focus on changes to the marginal cost schedule that may be revealed by the resolution of uncertainty and more generally by changes that may be precipitated by companion policies. Compared to existing adjustment mechanisms that take effect only for specific changes or are implemented with a lag, we find that price-responsive supply improves the effectiveness of trading programs by instantaneously adjusting the quantity of allowances supplied to the market based on directly observed market signals about compliance costs.

When companion policies reduce technology costs and thereby reduce the allowance market equilibrium price, a supply response in the allowance market channels the

---

6. Incremental actions involving a portfolio of policy approaches with the intent of affecting obstacles to carbon pricing have been described as policy sequencing (Meckling et al. 2015; Meckling et al. 2017; Pahle et al. 2018).

reduced costs to achieve further emissions reductions.[7] Such a decline in the quantity of allowances supplied supports the allowance market price. Alternatively, when allowance market prices rise to levels greater than what policy makers view as politically sustainable, a price-responsive supply schedule expands allowance supply to reduce the price, protecting the political stability of carbon pricing. All of this is accomplished according to fixed, pre-announced decision rules, reducing regulatory uncertainty, which strengthens investor confidence relative to delayed and unpredictable administrative adjustments. The rule-based adjustment allows market participants to evaluate likely future outcomes based on their private information about market circumstances.

The experiments show that price-responsive supply reduces price volatility in allowance market equilibria. In our experiments, we demonstrate an early and persistent effect on allowance price levels and variability. We also find evidence that reduced variability in prices will lead to reduced variability in auction revenues, at least with a linear price-responsive supply, which is important to the planning of investments to accelerate emissions reductions.

A price-responsive allowance supply introduces new policy parameters, including the number of price steps in the supply schedule, and their size, which together determine the average slope of the schedule. We use a realistic simulation model of CO_2 emission reductions in the US electricity sector to demonstrate that a greater number of price steps is expected to lead to less auction revenue variability, reinforcing the findings from the experimental setting. Further, we show in the simulation model that the price-responsive supply function leads to changes in investment and power generation that achieve additional emissions reductions compared to an inelastic allowance supply.

Price-responsive supply reduces the risks that firms and policy makers face due to the combined effects of price variability and regulatory uncertainty. In most situations price-responsive supply yields more stable carbon revenues, facilitating public sector investments that drive new technology, and attracting growing political support for carbon pricing through the development of coalitions that benefit from new economic opportunity. The automatic policy updating of price-responsive allowance supply can help policy makers achieve long-term emissions reduction goals with rules that have the attractive political properties of being both adaptable and durable.

In the next section, we review the development of the literature on prices versus quantities. In section 2, we discuss specifics about how price-responsive supply is implemented and provide theoretical support for its value in policy updating. Section 3 describes an experimental investigation of the performance of this new approach to allowance supply. Section 4 describes simulation modeling that is used to provide guidance in the design of

7. Indeed, it is virtually always the case that substantial revenues from allowance auctions are dedicated to functions, such as research and development and energy efficiency, that tend to lower the demand for allowances (Burtraw and Keyes 2018; Löfgren et al. 2018). These types of expenditures put an endogenous downward pressure on allowance prices.

price-responsive supply. Section 5 describes the policy experience with price-responsive supply. Section 6 provides a concluding argument that price-responsive supply should be a fundamental characteristic of allowance market design, and in environmental markets generally, and points toward a set of research questions that surface as a result.

1. LITERATURE AND HISTORY

The prices versus quantities debate stems from the 1950s literature on economic planning, which concerned whether an agency responsible for administering an economic plan within centrally planned economies or large firms should use quantity output controls or administered prices to implement the plan.[8] Weitzman's "Prices vs. Quantities" paper (1974) focused not on environmental policy but on the older debate over the central planner's task.[9] Weitzman asked whether, "for some isolated economic variable that needs to be regulated," is it better to administer the activity directly (set the quantity) or to fix transfer prices and allow profit-maximizing firms to decide what is produced. Weitzman showed that if the cost of production is uncertain, one needs information about the shapes of the cost and benefit schedules to minimize the deadweight loss arising from uncertainty over production costs. His main point was that neither prices nor quantities would be preferred for all situations. Adar and Griffin (1976) and Fishelson (1976) independently of each other and of Weitzman demonstrated the importance of the marginal cost and benefit functions in choosing between price and quantity instruments for emissions control.

At the same time, Roberts and Spence (1976) independently pointed out that, under uncertainty over abatement costs, the regulator generally can do better by mixing prices and quantities rather than depending on one or the other. They show that, if marginal pollution damages are constant, then there is no benefit to a mixed policy instrument. But, if marginal damages are rising, then a mixed system of "charges and licenses" that better approximates the social marginal damage schedule can outperform either instrument alone. The policy maker chooses the number of licenses to sell (the cap) and two "escape valve" prices—a price, p, at which more licenses will be made available if the emitter wishes to buy them and a subsidy value, s, granted to the firm for the extent to which its emissions fall below the number of licenses it owns. The policy maker sets the number of licenses so that, with expected abatement costs, the value of licenses will equal marginal abatement costs and will fall between p and s. If costs are unexpectedly high or low, then the price mechanism will be binding. With many firms, the license price can be established in a competitive market for the licenses.

8. Our focus here is only on the prices vs. quantities debate. For a detailed discussion of the history of the idea of pricing pollution, see Banzhaf (2020).

9. In the same year, Weitzman began a series of papers on common pool resources. This was the start of his long career of notable contributions to the environmental economics literature.

Various options for relaxing the fixed nature of the allowance supply have been included in climate policy proposals, including (1) a circuit breaker that would stall the rate of decline of an emissions cap at a specified price trigger, (2) an independent board established to manage the supply of allowances to keep prices within an acceptable range, and (3) a price collar setting a ceiling and floor on the price of emissions allowances (Pizer 2002; Aldy and Pizer 2009; Murray et al. 2009; Burtraw et al. 2010). In a variation on the last design, Karp and Traeger (2018) propose a redemption-value function that uses the market clearing allowance price to adjust the emissions value of each allowance certificate.[10] Existing programs do not have circuit breakers or independent boards, but they do have a regularly scheduled review that often leads to administrative adjustments to supply. A price collar of some form is also used in many existing cap-and-trade programs. At the price ceiling, additional allowances are made available, and at the price floor allowances are withheld from the market. In this framework, traders face a stepped supply curve for emissions allowances, which generally is a better approximation of the marginal damage function than either a fixed fee or a fixed cap.

Similar mechanisms are envisioned with respect to how an emissions fee could be adjusted to achieve a specific emissions goal, sometimes called an emissions assurance mechanism (Newell et al. 2005; Metcalf 2009; Aldy et al. 2017; Hafstead and Williams 2017).[11] Virtually all discussions start from one of the two conceptual approaches, quotas or fees, and explore modifications that make one look like the other or, potentially, offer a hybrid of the two approaches. The ability of the policy to respond to new information is key to performance in all these models. All these approaches differ in one important respect from virtually every commodity market—changes in allowance supply are implemented with a lag and some degree of discretion, while the supply of a commodity responds automatically and systematically with the equilibrium price obtained in the market.

Our approach is closest to Roberts and Spence (1976), who in an appendix generalized their hybrid mechanism by offering distinct assets at arbitrarily many steps (cap levels), each with its own price collar. Such a schedule can closely approximate any convex damage function, lowering the expected welfare cost of uncertainty over abatement costs. A challenge of implementing this multistep mechanism is that the markets for the increasing number of distinct assets, one for each step in the supply, become thin and will not be

10. The EU Emissions Trading System recently enacted a market stability reserve (MSR) to adjust the supply of allowances based on the number of allowances in circulation (the bank). The program removes allowances from auction and places them in the reserve when the privately held bank is larger than a specific size. Beginning in 2023, a portion of allowances held in the MSR is canceled when the reserve grows to a specific size. For analysis of some of the difficulties with this approach, see Holt and Shobe (2016a), Perino and Willner (2019), Rosendahl (2019), and Osorio et al. (2021).

11. See also the recent forum in a 2020 issue of the *Review of Environmental Economics and Policy*, https://academic.oup.com/reep/issue/14/1.

competitive.[12] To solve this dilemma, we propose a feasible implementation of the multistep market that does not divide up the allowance supply in a way that reduces competitiveness.

We also compare our approach to Pizer and Prest (2020), who provide analysis of the relative performance of emissions charges and emissions markets in a dynamic context where new information about emissions damages and abatement costs emerges over time and where policy can be updated in response. They describe a case where policy updating with error favors the use of emissions charges over emissions markets. In the next section, we will argue that price-responsive supply will be preferred to emissions charges, even in the case of erroneous policy updating. We also argue that policy makers may in fact make decisions intended to affect the costs and benefits of other policy choices.

2. THEORY AND APPLICATION

In this section, we describe price-responsive supply in more detail and present some theoretical results demonstrating its likely advantages over the traditional approach of emissions charges or fixed-quantity emissions markets. We will first show that, in the static policy context, where there is uncertainty over both marginal abatement costs and marginal damages, our approach will, under reasonable conditions, accomplish the improved outcome first suggested by Roberts and Spence (1976). We will then extend these results to the case of policy updating recently explored by Pizer and Prest (2020). Because price-responsive supply updates the cap level as new information about costs is revealed, it eliminates the advantage of emissions charges even in the case where the policy maker is subject to a randomly determined policy bias in later periods. Drawing on observed experience in existing emissions markets, we will also argue that responsive supply has a practical information advantage for policy makers who face the task of updating policy in light of new information about marginal abatement costs; that is, it is easier for policy makers to discover new information about marginal cost using responsive supply than using updated emissions charges.

2.1. The Static Case

Roberts and Spence prove that, for a given expected marginal damage schedule, a combined quantity and price regime reduces expected efficiency loss from mistakes in estimating the marginal abatement cost schedule. They demonstrate that dividing the cap into a sequence of assets each with its associated price-quantity step can approximate the

12. Roberts and Spence (1976, 204) show that "if one is prepared to introduce more than one kind of license, the penalty function can be made to approximate any convex damage function arbitrarily closely." But they also note that each of the asset markets must have enough participants "to ensure relatively competitive functioning."

marginal benefit schedule arbitrarily closely, eliminating the expected deadweight loss in the limit.

The proof posits a sequence of increasingly smaller markets for each set of licenses as the marginal damage schedule is divided into smaller segments. In principle, the resulting difficulty of market thinness can be overcome by recombining the separate assets into a common cap and selling the emissions allowances in an auction that determines both the quantity sold and the closing price based on the bids in the auction. In the appendix (available online), we provide a simple proof that selling a supply of allowances at auction along the expected marginal benefit schedule, implemented in an auction where bidders in the auction bid truthfully, eliminates the expected loss from mistakes in estimating marginal abatement costs, as is the case with the Roberts and Spence sequence of license markets.

Although truthful bidding is not generally guaranteed, there is reason to believe that feasible auction mechanisms may produce outcomes close to those designed for inducing truthful bidding. Montero (2008) proposed an auction mechanism that can induce truthful bidding when auctioning along a supply schedule. While Montero's suggested design would be difficult to implement in practice, Requate et al. (2019) show that for reasonably competitive auctions, little efficiency is lost by using a simple uniform price auction of the sort now used in many emissions trading programs across the globe.[13]

With the near-truthful bidding that can be expected in competitive uniform-price auctions, the deviation of the auction closing price from the marginal abatement cost can be expected to be small. If the allowance auctions can be made competitive, then the policy maker's actual policy choice is between the likely loss of efficiency in using a standard, uniform price auction to implement a price-responsive supply of emissions allowances and the likely efficiency loss from selecting the wrong emissions charge.

2.2. Policy Updating over Time

In a dynamic setting with policy updating in response to new information, we use the Pizer and Prest (2020) multiperiod model to show that, even with nonoptimal second-period policy making, price-responsive supply is superior to emissions charges under the assumption that the emissions allowances can be auctioned under conditions where the closing price at auction is equal to the marginal abatement cost. In Pizer and Prest, the information aggregation function of markets can go astray due to perverse policy outcomes in the second period. In effect, we do not want the market to anticipate bad future policies; the inability of market participants to act now on the perverse things the policy maker will do in the future is an advantage.[14] This result depends on the assumption that

---

13. Khezr and Mackenzie (2018a) discuss administratively manageable ways of implementing the Montero auction design.

14. We note that there are many variations on this theme, and one should be cautious about attaching too much generality to one of the infinite ways that policy making may diverge from

policy adjusts slowly compared to firm expectations. Firms observe new information about costs and benefits and act on this before observing the inefficient policy response. The price-responsive supply approach adjusts policy contemporaneously with the revelation of new information and, in doing so, eliminates the second-period cost that arises from firms acting before bad policies are revealed. So, even with the bad second-period policy, emissions charges no longer have the advantage.

Theorem: In the Pizer and Prest two-period case, where the policy maker uses a biased (shifted) estimate of the marginal damage function for policy updating in the second period, price-responsive supply is always better than an emissions charge.

Proof: Using the familiar setup from Weitzman (1974) where costs and benefits (in terms of emissions abatement rather than emissions) take the following forms:

$$C(q, \theta) = c_0 + (c_1 + \theta)(q - \hat{q}) + \frac{c_2}{2}(q - \hat{q})^2, \tag{1}$$

$$B(q, \eta) = b_0 + (b_1 + \eta)(q - \hat{q}) - \frac{b_2}{2}(q - \hat{q})^2, \tag{2}$$

where q is emissions abatement, $\hat{q}$ is the abatement target, η is uncertainty over damages, and θ represents uncertainty over costs, we know that it is optimal to set expected marginal benefits equal to expected marginal costs, $E[\mathrm{MB}(q, \eta)] = E[\mathrm{MC}(q, \theta)]$. But in the Pizer and Prest case, the policy maker is constrained to using a biased marginal benefit measure in the second period:

$$\mathrm{MB}(q, \eta) + \epsilon,$$

where ϵ is a measure of policy bias in period 2. Pizer and Prest compare an emissions charge to a cap. We replace their traditional quantity cap by a price-responsive auction in both periods. For each period, the emission cap (auctioned amount), which is the policy variable $\hat{q}$ in equations (1) and (2), is set so that price is equal to the marginal benefit of reducing emissions. In the first period, the auction is settled according to the expected marginal benefit curve:

$$p_1^{\text{auction}} = b_1 + b_2 \cdot q_1^{\text{auction}}.$$

In the second period, having observed η, we auction along the biased marginal benefit function used by the policy maker:

the economist's ideal. Presumably, if the bad policy came first, to be followed by good policy, then the calculus would tip in favor of emissions markets.

$$p_2^{\text{auction}} = b_1 + b_2 \cdot q_2^{\text{auction}} + \eta + \epsilon.$$

The price-responsive supply auction will be preferred whenever[15]

$$\frac{-b_2^2\sigma_\theta^2}{2c_2^2(b_2 + c_2)} < 0.$$

Since this is always true, then the price-responsive supply auction, even along the biased marginal damage function, is always preferred to updating a Pigouvian tax using the same biased marginal damage function.[16] So, a price-responsive auction outperforms a Pigouvian tax in all cases of policy updating studied by Pizer and Prest. QED

Once again, the comparison of the next best policy to price-responsive supply boils down to comparing any inefficiencies due to deviations in the auction price from the efficient price to any inefficiencies due to the setting of the incorrect sequence of emissions charges.

2.3. Applying Price-Responsive Supply to Policy Adjustment in Emissions Regulation

Modeling current carbon policies as dynamically equating marginal costs and benefits over a 200-year horizon is an important and useful academic exercise to guide research and development, investment, and policy design, but it is decidedly not how global warming policy is actually made in practice. As laid out by Baumol and Oates (1971), the practice of environmental policy making seems generally to start with a scientific result about damages due to the overuse of some resource or environmental service. The policy conversation often revolves around the question of how much of our limited regulatory budget should we be willing to spend to get the level of damages down to some acceptable level. The perceived cost of achieving given levels of reduction determines both the goal itself and the rate of approaching it. Policy makers establish a "budget" for spending resources on this good that we wish to "purchase" via our regulations designed to achieve the desired goal. For greenhouse gas emissions, this may amount to recursive rebalancing of emissions targets to equate marginal benefits and marginal costs in light of new information (Aldy et al. 2021) or to determining how quickly we can get to net zero (Stern and Stiglitz 2021) without busting the politically established social budget for doing so. This is not a one-shot decision. Everyone involved knows that new science will develop over time, production technologies will change, and new information about costs of achieving the goal will become available. The estimates of the marginal cost and marginal damage

15. Demonstration of this result is provided in the appendix.

16. The second-period result is the same for both policies. The superiority of the auction comes exclusively from the improvement in the first-period outcome.

schedules will be updated. That policies change with new evidence should not surprise anyone involved, but how they change can have a significant effect on outcomes.[17]

Existing emissions markets set emissions caps that fall gradually over time in order to help stay within the social budget by preventing the stranding of large amounts of existing capital assets and by allowing the development of new technologies that may lower the total costs of eliminating emissions. Initial prices in emissions markets, and carbon markets in particular, have often been lower than anticipated by policy makers and economists, and this has been accompanied by extensive banking of future allowances. The magnitude of the mistake in setting the cap level can be very large, as was the case, for example, for the European Union Emission Trading System (EU-ETS), the US Acid Rain Program, and the US Regional Greenhouse Gas Initiative (RGGI) cap.

In the case of the RGGI cap, between the time that the initial cap was set and the opening of the market, dramatic changes had taken place in related markets. In particular, the price of natural gas had fallen from nearly $14/million Btu (mmBtu) in 2005 when the initial memorandum of understanding launched the multistate initiative to $3 in 2009, the first year of the trading program. The RGGI allowance auction closed at the price floor for 11 consecutive auctions from the middle of 2010 until early 2013, when the RGGI states agreed to cancel allowances that had not been sold at the price floor and to halve the annual allowance allocation to accommodate a large private bank that had accumulated. A second adjustment to lower the cap occurred in the 2017 program review. Policy makers confided that the regulatory negotiations involving multiple jurisdictions in these decisions were "painful."

It also is likely that carbon policies are constrained away from what we might consider an efficient pathway by a variety of political and economic factors. Recent experience suggests that carbon markets operate in an environment driven to a substantial degree by companion regulations at various levels of government that direct technological change or that address complementary environmental goals such as local air and water quality or economic inequality (Fischer et al. 2021; Perino et al. 2021). Many of these companion policies would tend to suppress allowance prices under an existing cap. Policy makers seeking to achieve emissions reduction goals use companion policies to put a finger on the scales in favor of lowering emissions. These other policies can erode the price signal from the carbon market unless there is an explicit adjustment in the market to accommodate exogenous factors.

To strengthen the influence of a price signal and adjust regulatory stringency, carbon markets often apply administrative (discretionary) adjustments to the cap during periodic program reviews. For example, the RGGI memorandum of understanding incorporates a triennial program review to assess how the program is performing relative to previous

---

17. It is possible that a regime hostile to good governance may come into power. In this case, rational social planner models are of little help except possibly to estimate the costs of suboptimal policies.

expectations. Between reviews, the lag in regulatory response to the occurrence of unexpected prices perpetuates the earlier policy mistake. A lag between the realization of cost information and the policy response to that new information amplifies uncertainty about the likely response, undermines confidence in the market, and may fuel advocacy for additional sector-specific policies, pushing allowance prices even further from cost-minimizing levels, and potentially aggravating a vicious cycle by reducing the influence of prices in achieving the environmental goal (Flachsland et al. 2019). Setting the supply schedule in advance and then auctioning along this supply schedule provides both advance notice about future potential policy adjustments and instant feedback through the auction closing price about market expectations of future abatement costs. In fact, spot prices in secondary markets would likely signal firm expectations even before the auction occurred.

Alternatively, if we set price expecting to adjust it in response to the abatement costs inferred from the quantity response, the policy maker cannot simply calculate marginal abatement cost from the observed quantity outcome, because the quantity observed is a function of the error in estimating marginal costs, which the policy maker does not observe. The policy maker will need multiple observations on quantity to estimate marginal abatement costs. This will require a wait, possibly for years, to see how the induced investment in new capital stock has changed emissions. And by the time measurement is feasible, circumstances may have changed enough so that the estimates are not a good guide of current abatement costs. As we have already shown, eliminating policy lag reestablishes the superiority of using forward-looking market exchange rather than emissions charges for regulating emissions.

A price-responsive allowance supply provides a rule-based approach to adjusting stringency in response to realized information about abatement costs, and it adjusts cumulative stringency over time to automatically keep costs within the range of the agreed-upon regulatory budget. Periodic reviews remain essential because the policy maker's assessment of marginal damages (i.e., the allowance supply schedule) underlying the price-responsive supply will likely change over time; nonetheless, a price-responsive supply reduces the magnitude of an administrative adjustment that may be required during periodic program reviews.[18] A rule-based approach continuously responds to new information about allowance demand and marginal costs that is revealed in the allowance price, eliminating regulatory lag within compliance periods, boosting confidence in the market. In contrast, observed quantities in response to an emissions tax do not serve the same information aggregation function. Dynamic tax adjustments must be retroactive, while markets provide price information immediately and are forward looking.

---

18. An additional argument in favor of quantity targets is the desire to avoid possible existence of tipping points in damages. Price-responsive supply that is aligned with marginal damages accommodates this desire as the marginal damage curve becomes increasingly steep in the neighborhood of potential tipping points.

Figure 1 illustrates variations of price-responsive supply. Panel *A* illustrates an emission cap with a price floor and price ceiling; between these two price points changes in allowance demand map into changes in allowance price with no change in emissions. Panel *B* illustrates an additional price step above the price floor implemented with a reserve price for a portion of allowances offered in an auction, a design implemented as an emissions containment reserve in RGGI. If demand falls in this market, prices are greater and emissions are less than in the market represented in panel *A*. Panel *C* represents a multistep emissions containment reserve, which would be more continuously responsive to changes in allowance demand. Panel *D* illustrates that the regulator

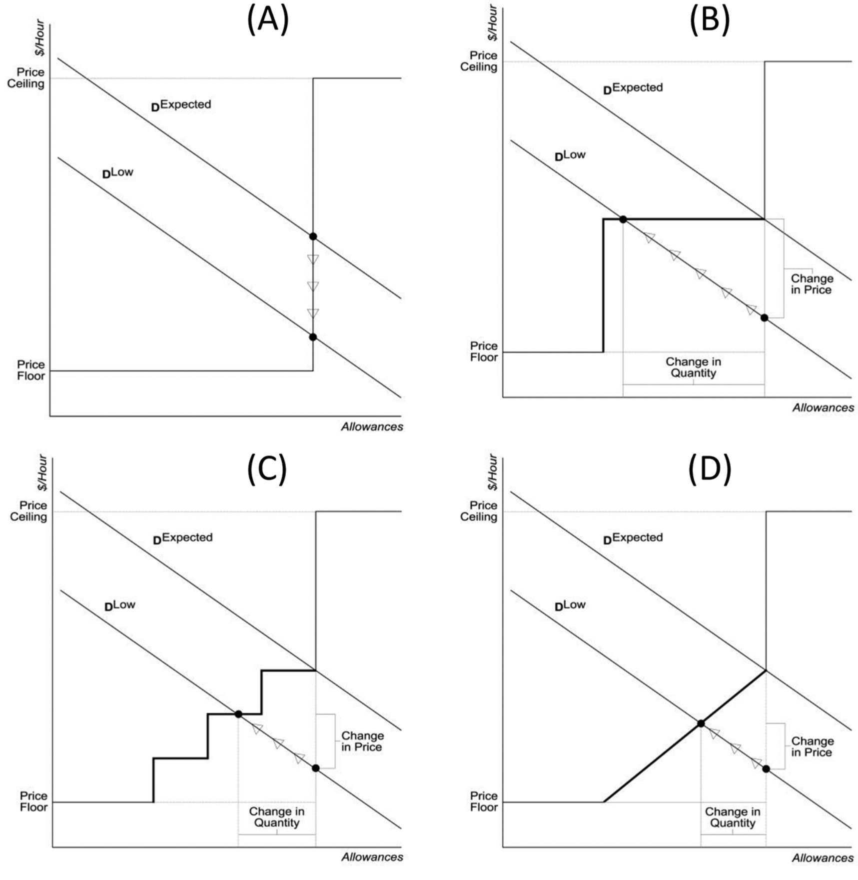

Figure 1. Variations of price-responsive supply in an emissions market. *A*, Emission cap with a hard price floor and soft price ceiling. *B*, Additional price step above the price floor implemented with a reserve price for a portion of allowances offered in an auction. *C*, A multi-step emissions containment reserve. *D*, Regulator specifies a smooth supply schedule and auction along this schedule.

could go so far as to specify a smooth supply schedule and auction along this schedule. In the next section, we will test the performance of the price-responsive supply mechanism in a set of experiments.

3. EVALUATING PRICE-RESPONSIVE ALLOWANCE SUPPLY IN A BEHAVIORAL CONTEXT

The implementation of a price-responsive supply uses bids in an auction to adjust the auction quantity. With added price steps, or with a smooth price-response function, the actual number of allowances sold will vary from the anticipated or nominal cap for that period. In our experiments, we test the performance of price-responsive supply that operates just below the expected, long-run equilibrium price given the nominal cap.[19] If the price-responsive supply mechanism in our setup is ever binding, then the allowance sales for that period will be below the nominal cap (vertical segment in the panels of fig. 1). This implies a tighter cap, compared to the same initial intended level without the price-responsive supply. As a result, forward-looking agents would anticipate the tighter cap and bid up allowance prices. This higher early price of allowances should persist over time, even in periods in which the price-responsive supply trigger point is not binding. Since a price-responsive supply mechanism dampens price reductions during slack demand periods, we expect price volatility to be reduced. Whether these expectations are realized depends on how well agents anticipate the effects of the mechanism and on the strategies they deploy when there is uncertainty about future payoffs from owning allowances. We test these implications using laboratory experiments with financially motivated human subjects.[20]

While there have been experiments testing the effects of reserve prices and price ceilings (Stranlund et al. 2014; Holt and Shobe 2016b), there have not been tests of how adding price-responsive supply affects market outcomes. In the case before us, we are interested in measuring the effect of adding price-responsiveness to supply in a simulated market designed to mimic some of the essential features of existing allowance markets.

Previous studies of the performance of price collar mechanisms, including its effect on banking, found that these mechanisms reduce the costs associated with uncertainty

19. The price-responsive supply mechanism can operate in either direction from the expected equilibrium price. We limit our focus to a mechanism operating just below the expected equilibrium price.

20. Experiments have been used previously to explore the effects of market designs in many of the key emissions markets implemented to date, including RGGI, the SO_2 allowance trading program, the eastern US NO_x market, the EU ETS, and the California CO_2 cap-and-trade program. In particular, laboratory experiments were a key component of the original consulting report proposing to set up the Regional Greenhouse Gas Initiative auction-based program (Holt et al. 2007).

over the future value of allowances (Cason et al. 2020).[21] Stranlund et al. (2014) used a multiperiod setting to show that, even in the presence of banking, price collars reduce price volatility, since market participants may not be able to smooth compliance costs sufficiently with banking alone. Cason et al. (2020) find that a hard price floor increases investment in lowering future abatement costs by increasing the expected return to that investment.

The results on the possible behavioral effects of price collars in static (single-period) settings is mixed. Perkis et al. (2016) found that soft collars were not as effective as hard collars in reducing price volatility. Friesen et al. (2019) examined different mechanisms for reducing high-price shocks. Their experiments suggest that supply reserves intended to dampen high-price shocks can under some circumstances lead to higher allowance prices.

Isaac and Plott (1981) reject the hypothesis that nonbinding price controls act as focal points for bidders in auctions for emissions allowances. More recently, Khezr and MacKenzie (2018b) identified conditions under which, in a static setting, an allowance reserve used to limit high allowance prices may alter auction clearing prices. In related experiments, Friesen et al. (2022) find that price collars may change bidding behavior in allowance auctions, pushing bids closer to price floors and ceilings.

In our experiments, we find, as expected, that price-responsive supply causes an early and persistent increase in allowance price and reduces price variability. We also find evidence that reduced variability in prices will lead to reduced variability in auction revenues, at least with a linear price-responsive supply function.

3.1. Experiment Setup

The experiments were implemented at the University of Virginia using the oTree experimental platform (Chen et al. 2016). Subjects were recruited in groups of 12 for sessions that lasted about 1.5 hours. Each session consisted of 30 periods or "rounds" that began with a 12-bidder auction and ended with the posting of earnings results based on the auction outcome and individual decisions on whether to use acquired allowances for current production or to bank them for future use. The auction mechanism used depended on the treatment; there were nine sessions in the no policy control treatment (as with fig. 1*A*). In addition, there were nine sessions with a single-step supply shift (fig. 1*B*) and nine sessions with a linear supply schedule (fig. 1C). In total, there were 27 sessions, each with a single fixed group of 12 bidders who competed in 30 consecutive auctions.

The laboratory setup presents subjects with a simplified version of a generic "permit" market, using neutral terminology that did not connect the permits needed for production to the notion of emissions allowances. Subjects can only acquire allowances in the auction; there is no spot market. The bidders interact through the determination

21. For a recent review of experiments on emissions allowance auctions generally, see MacKenzie (2022).

of the auction closing price based on all bids submitted, which in turn determines whether a supply adjustment is triggered.

Each of the 12 participants in a session controls four "capacity units" that produce one unit of output per period. Half of the participants were designated as "low users" (with low-emitting units), which require one permit per unit of output produced, while half were "high users" with units requiring two permits per unit of output. Output is sold at an exogenous price that was either $30 or $40 per unit, each with a probability of 50%. The shocks are uncorrelated. This output price variability induces periods with higher or lower current-use values for permits, which also depend on randomly determined costs that were independently generated for each bidder prior to the auction (details to follow). All sessions have the same random sequence of output prices and orthogonal cost realizations, so these factors are held constant across the three treatments.

Subjects are assigned a random cost of production for each of their capacity units in each round. Production costs vary uniformly on ($10, $28) per unit of output for low emitting units and on ($1, $28) for high emitters. The net value of a permit for current use is the difference between the current product price and the realized production cost for the unit, divided by the required number of permits needed to operate the capacity unit.[22]

Each of the 30 periods in a session begins with a uniform-price auction of the available allowances, sold at the highest rejected bid, and a production stage, in which subjects choose how many of their capacity units to run. The running of capacity units determines the number of permits that must be retired for compliance with the permit requirement. Running a unit without owning a permit to cover its operations incurs a substantial penalty ($35 per missing permit).

The number of allowances sold at auction starts at 66 in the first period of a session and declines by one allowance each period, going down to 37 permits auctioned in period 30. This pre-announced tightening of the cap provides an incentive to bank permits early on in anticipation of increased scarcity later. Previous experiments have shown participants to be adept at smoothing the supply of allowances over time, although they do so imperfectly (Shobe et al. 2014). This implies that the price in early sessions should provide a good signal about the long-range tightness of the cap. If there were no smoothing, we would expect to see the price rise as the cap falls, but with effective smoothing, the price in early periods will be very similar to the price in later periods.[23]

22. In low-price periods, this net value ranges from $1 to $15.50 for high emitters and from $2 to $20 for low emitters. In high-price periods, this net value ranges from $6 to $19.50 for high emitters and from $12 to $30 for low emitters.

23. For simplicity, the experiments are structured to have a zero discount rate, which yields a flat predicted allowance price trajectory in a model with perfect foresight. This simplification does not change the key results.

All auctions have a reserve price of $5, with no bids accepted below that price. The single-step supply includes 16 allowances at the $8 step. In contrast, the linear ramp supply declines smoothly from the trigger price of $8 down to the auction reserve price of $5. All sessions also include a "cost containment reserve" (soft price ceiling) of 10 allowances that will only be released into the auction if it would otherwise close at a price above $12. Therefore, all sessions have precisely the same structure except for the introduction of a price-responsive supply and the way it is implemented. The appendix contains instructions that participants received about changes in the supply of permits. The only difference in the instructions between the three treatments is description of the supply response on page 3 of the instructions.

Subjects were recruited from the University of Virginia student body. There were 324 participants in total, each of which received a $6 participation payment, plus a payment of 2.5% of experiment earnings that were added to an initial cash endowment of $150 lab dollars. Total cash earnings averaged about $30 per subject.

3.2. Experiment Results

Since there is no discounting in this setup, optimal permit management implies that permits will be used whenever their current use values are above a cutoff. To determine this cutoff, the total cap S (the sum of all auction quantities in all 30 periods without any supply withdrawals) is used to determine the S highest current use values (price minus cost divided by the required number of permits for that unit). The long-run supply (total cap) and demand (based on valuations in all periods) are equal at a price of $8. This would be the price if all allowances were used optimally, that is, if there were perfect smoothing of cost and price effects across periods, that is, banking of permits in early periods with current use values below this cutoff. If some allowances will be retired, then the increased scarcity of allowances would imply a higher long-run allowance price than the $8 level. Of course, subjects do not have the information about random demand realizations or the distributions of random production costs necessary to determine the optimal $8 current-use cutoff. If subjects anticipate a supply adjustment mechanism to be binding in some periods, then they might also expect that the aggregate supply of allowances will be reduced.[24] Given the ability of people to anticipate future scarcity, the likelihood that a supply reduction will be triggered will reduce the long-term supply of allowances, which should raise prices in initial periods relative to prices in the no supply adjustment control treatment.

Figure 2 clearly shows a pattern of higher average allowance prices for sessions with a responsive supply than for control treatment sessions. On average, prices for both types of responsive supply, step and linear, are higher than for the no-responsiveness

---

24. Salant et al. (2022) describe and test experimentally a storage model with a nonbinding price floor. They demonstrate the mechanism results in price floors pushing prices upward even when they are not binding.

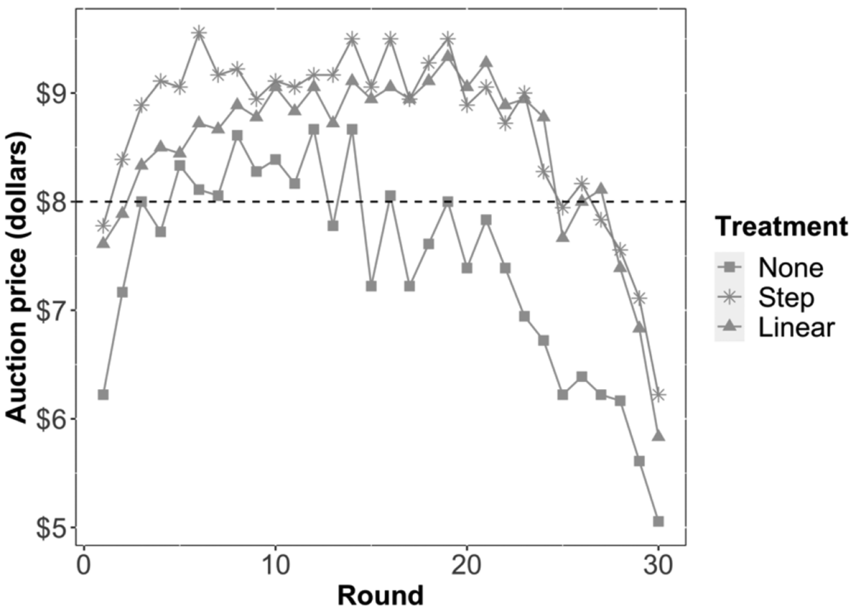

Figure 2. Average auction price by treatment by round

case in all periods. This difference is statistically significant, although the small difference between prices in the step and linear treatments is not.

> **Experimental result 1**: Allowance prices are higher in price-responsive supply treatments.

Support: In order to avoid end effects, the average price for each session was calculated for the first 25 periods. With nine average price observations per treatment, the usual assumptions needed to justify a standard *t*-test are questionable. Therefore, we used a nonparametric permutation test that determines the proportion of the 48,620 permutations of treatment labels that produce a difference in treatment average prices that is as large as or larger than the observed difference in either direction (a two-tailed test). In this manner, the null hypothesis of no effect of the step supply relative to no supply response is rejected ($p = .0054$). Similarly, the null hypothesis for the linear supply relative to no supply adjustment is rejected ($p = .022$). These and all other permutation tests to be reported are two-tailed tests. These p-values for tests are almost identical to those obtained from a rank-based nonparametric test (Mann-Whitney). The numerical approach, using permutations of actual prices, can be justified when the underlying data are not ordinal.

Because the quantity of permits is not fixed, but rather is immediately responsive to demand during the auction, we expected to observe lower price volatility with a price-responsive supply. The lower price volatility of the responsive supply treatments relative to the no-responsiveness control is apparent from the sawtooth patterns in figure 2, which leads to our second result:

Experimental result 2: Allowance price volatility is lower in price-responsive supply treatments.

Support: Using the first 25 rounds of auction price data, the average of the absolute values of round-to-round auction price changes is calculated for each session, which provides nine independent variability measures per treatment. The average absolute value of the auction price change is 0.79 for the no-responsiveness sessions and 0.41 for the linear sessions. Only 1.4% of the permutations of the treatment labels yield a variability difference as large as this in either direction, so the result of a two-tailed permutation test of the null hypothesis of no difference can be rejected ($p = .014$). Similarly, for the step treatment, the average absolute price change is 0.44, which is also significantly different from the control treatment ($p = .038$). Finally, the null hypothesis that auction price variability is the same for the step and linear supply cannot be rejected ($p = .74$).

The use of a price-responsive supply could also change the incentive to bank allowances for the future. In theory, early banking could go either up or down in response to the presence of the supply adjustment mechanism. The lower observed price variability in the responsive supply treatments (result 2) suggests that there is less of an opportunity to accumulate permits at bargain prices in early low-demand periods, and as a result, banking could be reduced. Working against this possibility, some participants may wish to hold a higher bank in anticipation of fewer permits being available in the future. Our results are suggestive that the increased cost of accumulating a bank can be a powerful counterbalance to the desire to hold a larger bank in anticipation of future scarcity. As seen in figure 3, the average bank tended to be lower in the responsive supply treatments, especially in early periods.

Experimental result 3: The evidence that allowance banks are lower in price-responsive supply treatments is mixed but strongest in the early periods.

Support: The test uses the total bank for each session, averaged over the early periods (1–15). The null hypothesis of no difference is not rejected for step treatment ($p = .14$, two tails), but the analogous null hypothesis of no effect is rejected and for the linear treatment ($p = .098$, two tails). Using periods 1–25, the bank for the step treatment is not significantly lower than the no-response treatment ($p = .248$), but it is significantly lower for the linear treatment ($p = .094$).

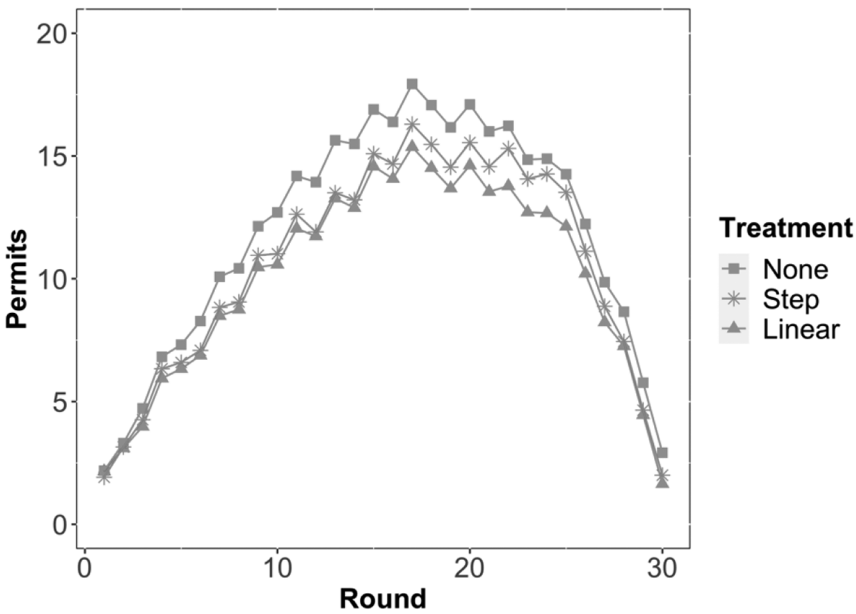

Figure 3. Total banked allowances by treatment by round

There is good reason to believe that lower price variability is generally beneficial. Policy makers also care about the revenue stability because revenue variability complicates budget planning. A price-responsive supply of allowances should tend to dampen price variability compared to a fixed cap. Some of the variation that would be loaded on price is now shared with variation in quantity, so the net effect on auction revenue is unclear. Figure 4 shows average auction revenues in each round by treatment. The largest revenue changes are for the middle periods of the no-response control treatment, as indicated by the sawtooth pattern of the lower line for that treatment. Of course, the significance of the result depends on the pattern of variability measures across independent sessions of each treatment, not just on overall averages. It turns out that auction revenue variability is significantly lower for the case of a linear supply, but not for a stepped supply, which leads to our fourth result:

Experimental result 4: Revenue volatility is lower with the linear supply but not with a step.

Support: Using the first 25 rounds of auction revenue data, the average of the absolute values of round-to-round revenue changes is calculated for each session. The treatment average auction revenue change is 46.37 for the no-response sessions and 28.37 for the linear treatment. Only 2.1% of the permutations of the treatment labels

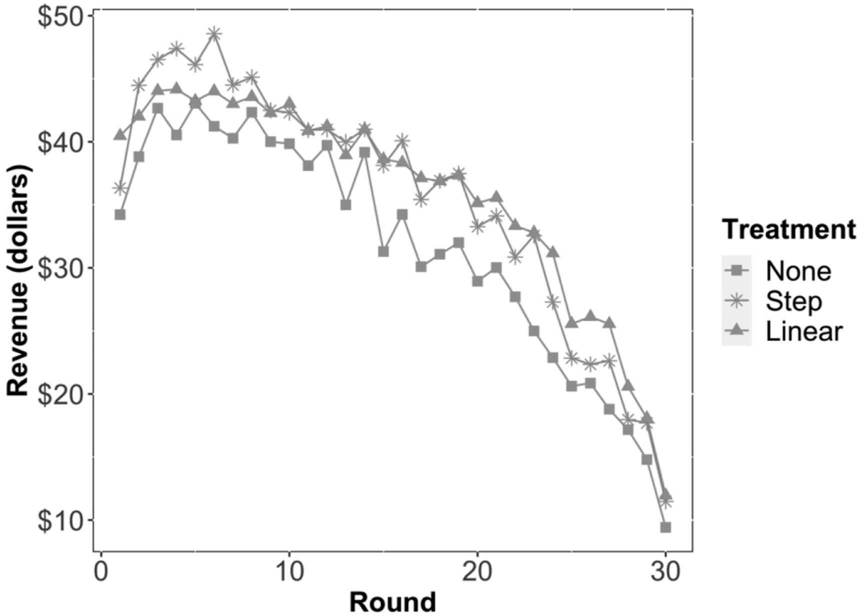

Figure 4. Total revenue per subject by treatment by round

yield a difference as large as this in either direction, so the result of a two-tailed permutation test is $p = .021$. In contrast, for the step treatment, the average absolute revenue change is 40.33, which is close to the observed 46.37 average for the control sessions, and the difference is not significant ($p = .48$, two-tailed). Finally, the null hypothesis that auction revenue variability is the same for the step and the linear responses can be rejected ($p = .056$).

Table 1 provides a summary of the treatment averages. As expected, we find that the price-responsive supply, step or linear supply, causes an early and persistent increase in auction prices above the $8 level (second column) and that price variability is reduced (third column). The price-responsive supply of allowances also results in a reduction in the size of the bank held during the early periods, although the difference is not uniformly significant. The reduced bank may be, in part, the result of reduced price variability due to the dampening effect of the price responsive supply. Finally, as shown in the right-hand column of table 1, the price responsive supply (step and linear) resulted in reduced revenue variability for the linear supply.

The final issue is how the different treatments affected the total production of both low and high users combined. The expected effects of the policy regimes on output in the product market are ambiguous, with two channels of influence, which operate in opposite directions. By canceling some allowances, and hence reducing aggregate

Table 1. Experiment Treatment Averages

Treatment	Average Price[a]	Price Variability[a]	Average Bank[b]	Revenue Variability[a]
No response	7.71	.79	10.27	46.37
Stepped supply	8.97***	.44**	9.04	40.33
Linear supply	8.74**	.41**	8.75*	28.37**

Note. The average price for each session is the average over the first 25 auction prices. Price variability is the average of the absolute values of period-to-period auction price changes in rounds 1–25. Revenue variability is calculated similarly using averages of absolute values of auction revenue changes in rounds 1–25. Average bank for a session is the average of the total bank held by all subjects in a session over the first half, rounds 1–15.

[a] Indicates an average for the first 25 periods.

[b] Indicates an average for the first 15 periods.

* $p < .10$ (for two-tailed permutation tests vs. no PRS).

** $p < .05$ (for two-tailed permutation tests vs. no PRS).

*** $p < .01$ (for two-tailed permutation tests vs. no PRS).

supply, the step and linear treatments would tend to raise the cost of production and reduce output. But the increased price due to the lower long-run supply would also cause a shift of production from high emitters to low emitters, a shift that would increase production for a given level of emissions. Our theory does not predict which of these effects is likely to be larger, so we are not able to say whether the amount of final output should rise or fall in our setup. In our experimental sessions, production of the final output is not statistically different among the treatments.

4. SIMULATIONS

The conceptual discussion in section 2 demonstrated the efficiency advantage of aligning the price-responsive supply of allowances with marginal damages. We use climate policy as a specific case for simulating price-responsive allowance supply, although the concept has general applicability in emissions markets. Some authors have suggested that the marginal damages of greenhouse gases are constant, implying that a tax would be preferred on efficiency grounds. But policy makers are likely to recognize air quality cobenefits that confer slope to the marginal benefits of carbon emissions reductions. And in the long run, the marginal benefit of greenhouse gas emission reductions is not constant, while indeed marginal abatement cost may become constant due to emergence of backstop technologies, suggesting a possible preference for quantity instruments in the Weitzman context. Practically speaking, having established that there are efficiency advantages in moving away from inelastic allowance supply in emissions markets (a strict cap), policy makers are left with the challenges of knowing the proper value of marginal damages for setting policy and implementing prices that approach this value. As policy makers incrementally expand carbon policies, one general concern is the volatility of

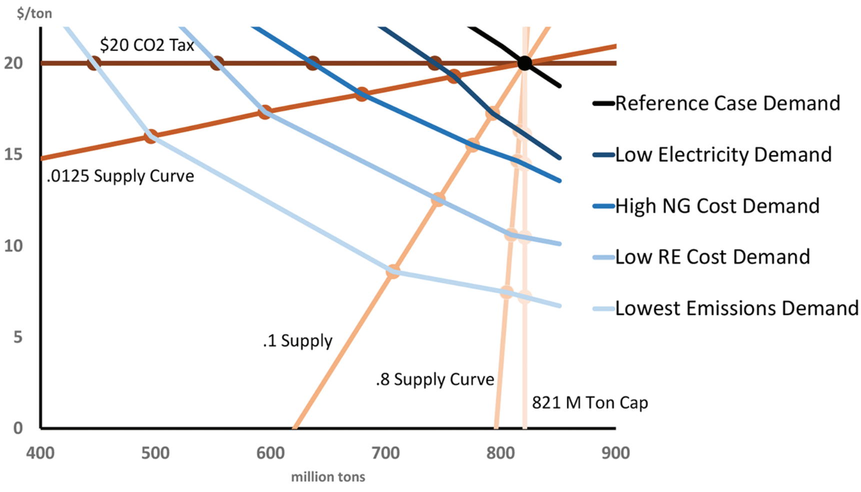

Figure 5. Electricity market equilibria with different supply and demand schedules

revenue in allowance auctions because revenue is tied closely to investments that are intrinsic to the design of climate policies.

To examine approaches to militate against revenue volatility and provide practical advice for policy makers we use a dynamic linear programming simulation model of the US electricity sector. The model solves for investment in and retirement of generation capacity over a 23-year horizon, with annual operation of the electricity system represented in four time blocks in each of three seasons. Electricity market equilibria are solved at the state level, allowing state-level representations of environmental policies and regulatory practice, with interstate transmission capability calibrated to observed transactions. We focus on 2026, a year sufficiently in the future to accommodate comprehensive carbon pricing in the electricity sector.[25]

We calibrate the CO_2 emissions market outcomes with an inelastic national electricity sector emissions cap of 821 million short tons, which is achieved at a price of $20/ton under reference case assumptions (fig. 5). We then explore alternative scenarios in the electricity market that reduce emissions allowance demand in the carbon market. These alternative scenarios include a 4% annual decrease in renewable cost, a 1% annual decrease in electricity demand, and a 20% increase in natural gas fuel cost in 2026. With inelastic emissions allowance supply, changes in the derived demand for allowances map entirely into changes in the price of allowances. Low renewable costs yield an allowance price of $10.47, low electricity demand yields an allowance price of $16.10, and the

25. The model does not allow intertemporal banking of emissions allowances to facilitate comparison with an emissions tax.

high natural gas price yields an allowance price of $14.46.[26] If all three of these outcomes occur simultaneously, we observe the lowest emissions demand and a price of $7.20. Conversely, allowance prices would increase if scenarios caused allowance demand to increase above reference case levels. Here we focus exclusively on outcomes that lower allowance demand, with a special interest in outcomes regarding renewable prices and energy demand that are likely to be affected by coincident policies. Cases when allowance demand falls in relation to reference case projections align with the dominant observed phenomena in environmental markets as discussed in section 2. An inelastic supply of emissions allowances translates all changes in market fundamentals into price movements, yielding high price and revenue variability. Total allowance revenues fall from $16.41 billion in the reference case to $5.91 billion when all three outcomes occur simultaneously in the lowest demand scenario. Other things being equal, policy makers would prefer a design that lessens revenue variability,[27] an important concern for emissions trading programs that increasingly are described as "cap and invest" meaning that emissions reductions are explicitly intended to result from the investment of auction revenue.

We explore the effect of price-responsive supply on revenue volatility. Two main new features characterize the design of a price-responsive allowance supply schedule: the (average) slope of the schedule and, if it has discrete steps, the number of steps. In the simulations displayed in figure 5, we consider three continuously differentiable linear allowance supply schedules (ramps with no discrete steps) with slopes of 0.8, 0.1, and 0.0125, in the southeast quadrant illustrating the region where demand for allowances falls relative to the reference case. Across potential equilibria identified by intersections of allowance supply and demand, less slope in allowance supply (more elastic supply) results in greater responsiveness in the emissions outcome when the demand for allowances falls compared to the reference case. Concurrently, the less slope in allowance supply, the less responsive is the allowance price.

Among the equilibria illustrated along each demand curve, we calculate revenue to be greater on the flatter allowance supply curves than on the steeper ones, although not as great as under the carbon tax which is perfectly elastic.[28] In general, whether revenues are best preserved under a quantity or price constraint depends entirely on the price elasticity of demand in the allowance market.[29] At the reference case equilibrium, we calculate an

26. Allowance prices fall with higher natural gas prices because in this neighborhood of carbon prices the main margin for achieving incremental emissions reductions is substitution between gas and renewables. Little coal generation remains in the system.

27. For example, in May 2021 Washington enacted its Climate Commitment Act to create an economy-wide cap on greenhouse gas emissions, with a central feature of an allowance auction that would yield revenue for investments.

28. We assume under every scenario that unsold allowances are permanently canceled and do not enter the market at a later date.

29. If demand falls, the change in revenue would be minimized if the slope of the supply schedule were the negative inverse of the slope of the demand curve.

allowance demand elasticity of –0.6, and at every allowance market equilibrium over the range of carbon market outcomes we examine we find allowance demand to be inelastic. With such a low elasticity of demand, revenues are best preserved under a CO_2 tax equal to $20, the reference case price observed in the carbon market. Under a tax, low electricity demand results in a reduction in emissions from 821 million tons to 743. Low renewable costs result in emissions of 554 million tons, and high natural gas costs result in 637 million tons. All three changes together in the lowest demand scenario result in emissions of 447 million tons. In this case, revenues from the carbon price fall in proportion to the change in emissions, from $16.41 billion in the reference case to $8.93 billion.

Revenue stability is one of many criteria shaping program design and other criteria including aversion to potential emissions increases often tip toward an allowance trading approach. The general guidance in an allowance-based system in the US electricity sector is that if demand for allowances falls, revenue will be more stable the flatter is the supply schedule.

The second new design feature relevant to the implementation of price-responsive allowance supply is the number of discrete price steps.[30] To enable comparability of price steps with the price ramps, we construct step function supply curves with an average slope of 0.1, the middle case illustrated in figure 5, for prices above the price floor of $9.50. We assume that steps along the supply schedule are of even height and evenly spaced.

Taking the reference case demand for emissions allowances from the simulations as a point of departure, we build a discrete distribution of 14 parallel and evenly spaced derived demand curves for emissions allowances ranging from the expected value in the reference case (with an equilibrium price of $20/ton and emissions of 821 million tons identified previously) to the lowest derived demand that intersects the supply curve where it hits the price floor.[31] With this distribution of potential realizations of allowance demand, we compare a linear supply curve (ramp) and a step function with the same average slope and up to nine price steps. The demand curves and three supply schedules including one step, three steps, and a ramp are illustrated in figure 6*A*.[32]

In the figure, one can observe that for any specific realization of demand, auction proceeds might be greater under the one step, three step, or ramp supply schedule. However, under restrictive but intuitive conditions about the distribution of expected changes in allowance demand, we find that a greater number of price steps maps into greater expected revenue on average over the indicated range of possible realizations with the

30. An additional feature is how the steps adjust over years, which influences investment. The price floor in the Western Climate Initiative and the emissions containment reserve in RGGI increase at 5% per year (real dollars).

31. The reference demand curve has an elasticity of –0.1 at the $20 allowance price, which establishes the slope of all the demand curves.

32. If the allowance supply schedules were illustrated in fig. 6 as continuing above the reference point, the reference point would bisect a vertical portion of the steps.

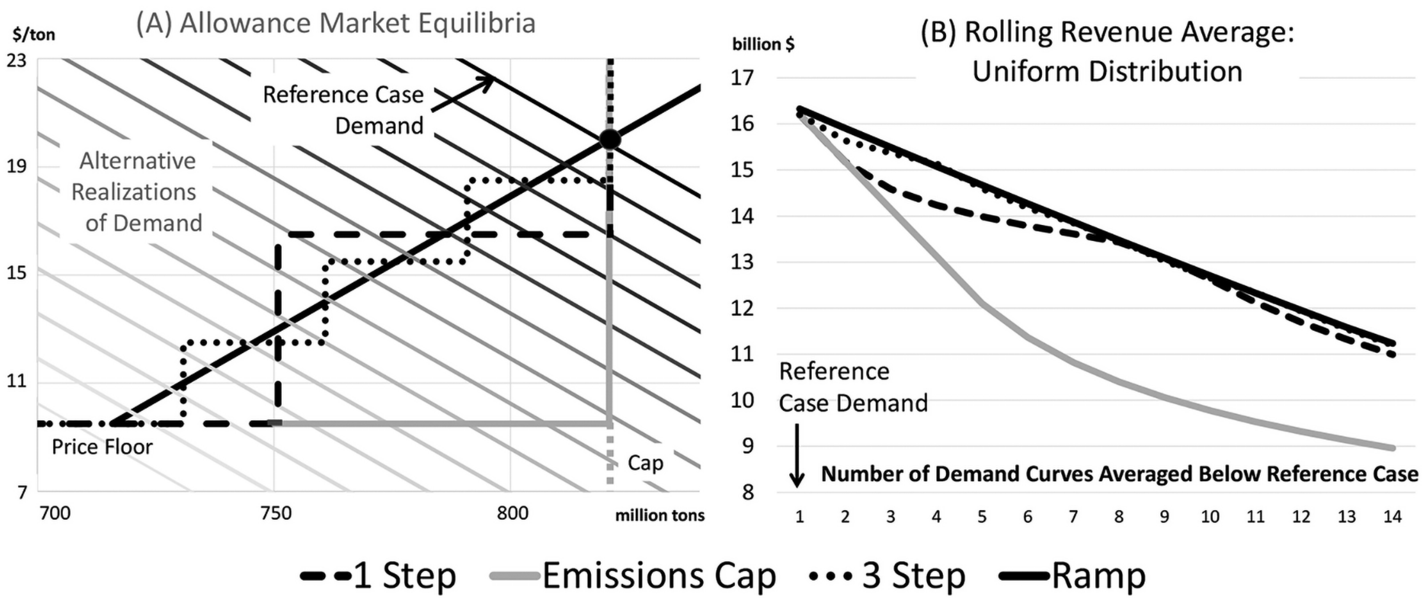

Figure 6. Realizations of allowance demand and alternative supply schedules

linear schedule implicitly having the most steps. We calculate expected revenue initially assuming a uniform probability distribution of potential demand outcomes anchored at the reference case demand with a support spanning to include from 1 to 14 of the demand curves illustrated in figure 6*A*. Variation in demand for allowances stems from the policy and secular shocks to costs like those modeled above. Figure 6*B* shows a rolling average of revenue as the support for possible demand realizations is expanded, beginning with the reference case demand curve, and broadening to include additional lower demand curves in sequence all the way down to where demand intercepts the price floor. For up to nine price steps and a continuous ramp, revenue is always greater at each potential equilibrium than under an inelastic allowance supply (a cap), and multiple price steps always outperform a single price step.[33] We also examine a one-sided triangular distribution of demand curves with the mode at the reference case and varying support ranging from one to all 14 demand curves, we examine supply curves with steeper (0.8) and flatter (0.0125) slopes illustrated in figure 5, and we vary the slopes of the parallel demand curves including steeper (−0.2) flatter (−0.05) demand, and find the same result.

In summary, in the electricity sector, which has been most often included in emissions markets, and where the demand for emissions allowances is generally inelastic, we conclude that revenue will be closer to expected revenue if alternative realizations of uncertain allowance demand results in relatively greater changes in the quantity of emissions and smaller changes in the price. This argues for a supply schedule that has a relatively flat slope. For a supply schedule with a given average slope, the greater the number of steps, the greater the resilience of allowance revenue to potential changes in allowance demand.

5. POLICY EXPERIENCE

The first application of price-responsive allowance supply was the adoption of an auction reserve price floor at the 2009 launch of the Regional Greenhouse Gas Initiative, which currently covers the electricity sector in 12 northeastern US states. RGGI also offers a quantity of allowances in addition to the anticipated emissions cap at a cost containment reserve price step. The Western Climate Initiative CO_2 economy-wide trading program comprising California and Quebec also has a price floor and offers allowances in addition to the anticipated cap in a two-step allowance price containment reserve. In addition, it has a price ceiling at which an unlimited supply of compliance instruments (i.e., allowances) become available. In both these programs, the price floor has been periodically triggered automatically constraining the number of allowances that are sold in the auction, and in both programs the price has also risen off the price floor in subsequent auctions. In RGGI, the cost containment reserve has been triggered twice.

In addition to these automatic adjustment mechanisms, both programs have administered reductions in allowance supply during regularly scheduled program reviews. This

33. Under this discrete formulation, the differences in revenue vary inconsistently and to a diminishing degree when comparing supply curves with greater than one step.

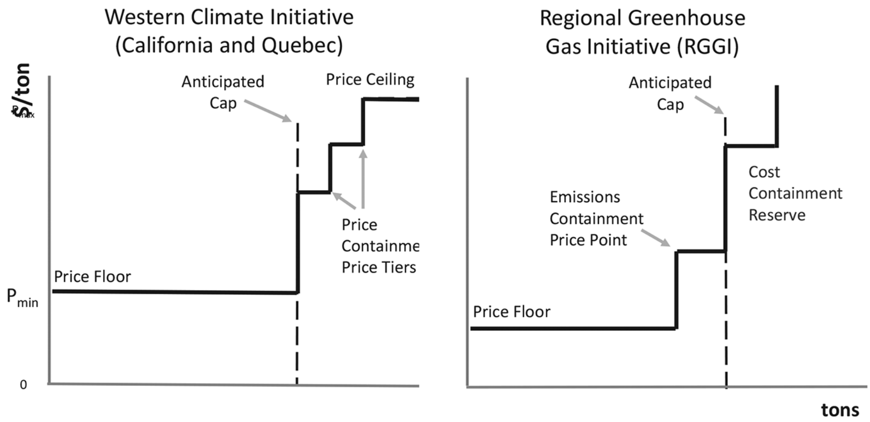

Figure 7. Actual price-responsive supply schedules

approach has been especially important in RGGI, where in 2012 the price had rested on the price floor for 11 consecutive auctions coincident with the accumulation of a large bank of unused allowances. When a large bank accumulated again three years later, RGGI again reduced allowance supply.

These programs originally unfolded in the face of several uncertainties that primarily resolved to depress emissions and the demand for allowances. The expansion of natural gas supply has contributed to reduced coal-fired generation. The economic recession in 2008–9 reduced electricity demand, which remained low as the economy recovered. State and federal policies to support renewable technologies put downward pressure on emissions allowance prices, as have programs to promote energy efficiency in buildings and vehicles. A large portion of the auction proceeds in both trading programs is explicitly directed toward strategic energy investments to accelerate emissions reductions.

RGGI has been especially affected by the policy ambitions for emissions reductions of many of its member states that collectively exceed the regional goal.[34] To avoid a waterbed effect that might result from state-level ambitions, RGGI adopted a new automatic adjustment feature, known as the emissions containment reserve, to provide additional price responsiveness in the allowance supply curve taking effect in 2021, which applies a reserve

34. Modeling exercises for RGGI in 2016 using the Integrated Planning Model (IPM) anticipated a price path rising from $7 to $9 between 2020 and 2026 (https://www.rggi.org/program-overview-and-design/design-archive/2016-materials). By April 2017, due to changes natural gas price forecasts, updated projections for electricity demand and the cost and performance of renewables, and anticipation of state policies and additional renewable imports from Canada, IPM's projected allowance price for 2020 and the subsequent decade fell to just above $2.

price above the price floor level to 10% of the allowances offered in the auction.[35] Consequently, we observe that the two North American carbon markets show an evolving transition in emissions trading away from inelastic allowance supply that has framed the debate over prices versus quantities for nearly 50 years, toward increasing price responsiveness of allowance supply. A schematic of the allowance supply schedules in these programs is shown in figure 7.

6. CONCLUSION

A key challenge for policy makers in implementing emissions pricing schemes is how best to introduce policy flexibility and to update regulatory stringency as new information arises about marginal abatement costs, broadly defined. We propose a practical strategy for implementing a mechanism first proposed by Roberts and Spence (1976), which effectively eliminates the distinction between price and quantity instruments for controlling emissions. Roberts and Spence proposed a sequence of markets, each with its own price collar. Until now, this strategy has not been thought practical, since each of the small asset markets would be too thin for competitive price discovery. Our strategy, price-responsive supply, determines price and quantity simultaneously by using the policy maker's best estimate of the marginal damage schedule as the supply function that is used to settle allowance auctions. We show theoretically that this mechanism, if implemented in a competitive auction, dominates a pure price mechanism even in cases that previously have appeared to favor a price mechanism over emission trading.

The key advantage of price-responsive supply is that it automatically adjusts the supply of allowances contemporaneous with the formation of the market price and, hence, contemporaneous with the discovery of new information by market participants. A price-responsive allowance supply supports a dynamic process that enables policy makers to use incentive-based regulation to achieve environmental quality goals with regulations that interact positively with companion efforts to achieve multiple, related policy objectives. Price-responsive supply also helps overcome the waterbed effect and establishes a predictable framework for policy updating that can reduce regulatory risk.

We have argued here that, with a price-responsive allowance supply that automatically adjusts to new information about abatement costs, emissions markets are preferable to a feasible emissions charge, including a charge that is updated over time. Our experiments demonstrate that price-responsive supply can be implemented through the auctions currently used in existing emissions markets and that it can reduce price and revenue volatility. Both RGGI and the Western Climate Initiative have made initial steps toward implementing elements of price-responsive supply, which have been important to the success of these cap-and-trade programs.

35. Washington (state), which passed legislation in 2021 to enact an economy-wide CO_2 trading program beginning in 2023, also includes an emissions containment reserve.

We test the price-responsive supply mechanism in the laboratory setting. In multiround sessions, during each round subjects purchased allowances in a uniform price, sealed bid auction and then banked allowances or used them for current compliance. We found that the presence of the price-responsive supply schedule as the auction settlement rule has an early and persistent effect on allowance price and reduces price variability. We also found evidence that reduced variability in prices led to reduced variability in auction revenues, at least with a linear price-responsive supply schedule. We find suggestive evidence of reduced banking of allowances, at least in early periods.

Even if policy makers embrace the efficiency advantage of moving away from inelastic (fixed) allowance supply, they face a challenge in identifying the marginal damage schedule that would enable them to set an efficient supply schedule. They are faced, nonetheless, with the practical challenge of implementing emissions pricing and a specific challenge in the variability of auction proceeds, which are important to program-related investments. We explore the performance of price-responsive allowance supply in simulation modeling of CO_2 emissions reduction to develop policy guidance. Where demand for emissions allowances is inelastic, we find that a flatter allowance supply schedule has less revenue variability when demand for allowances falls below expected levels. For any given slope in an allowance supply schedule, we find that more price steps will reduce revenue variability.

Our results have particular relevance for efforts to control global warming. Climate policy is taking shape in international, national, and subnational contexts where the balancing of economic and political costs and benefits is part of a regulatory negotiation within and between jurisdictions. National and subnational environmental goals are derived and evolve in a noncooperative coordination setting with substantial uncertainty. Most jurisdictions that have adopted carbon pricing have chosen a cap-and-trade approach, with emissions targets that are clearly not globally optimal and are not static. The inherent flexibility of the price-responsive supply mechanism may facilitate the wider adoption of carbon pricing globally.

Price-responsive supply contributes to a resolution of the difficult choice between price versus quantity instruments that has characterized nearly 50 years of economic debate. Such a schedule moves toward a design for environmental markets that more closely resembles that of other commodities. We anticipate that price-responsive supply can help reduce the costs of administrative adjustments to program stringency, making the policy more durable and helping to increase the influence of emissions pricing in driving emissions reductions moving forward.

REFERENCES

Adar, Zvi, and James M. Griffin. 1976. Uncertainty and the choice of pollution control instruments. *Journal of Environmental Economics and Management* 3 (3): 178–88. https://doi.org/10.1016/0095-0696(76)90017-6.

Aldy, Joseph E., Marc Hafstead, Gilbert E. Metcalf, Brian C. Murray, William A. Pizer, Christina Reichert, and Roberton C. Williams III. 2017. Resolving the inherent uncertainty of carbon taxes: Introduction. *Harvard Environmental Law Forum* 4:1–13.

Aldy, Joseph E., Matthew J. Kotchen, Robert N. Stavins, and James H. Stock. 2021. Keep climate policy focused on the social cost of carbon. *Science* 373 (6557): 850–52.

Aldy, Joseph E., and William A. Pizer. 2009. Issues in designing U.S. climate change policy. *Energy Journal* 30 (3): 179–210.

Banzhaf, H. Spencer. 2020. A history of pricing pollution (or, why Pigouvian taxes are not necessarily Pigouvian). Working paper 27683, National Bureau of Economic Research, Cambridge, MA. https://doi.org/10.3386/w27683.

Baumol, William J., and Wallace E. Oates. 1971. The use of standards and prices for protection of the environment. In *The economics of environment*, 53–65. Cambridge: Springer.

———. 1988. *The theory of environmental policy*. 2nd ed. New York: Cambridge University Press.

Burtraw, Dallas, and Amelia Keyes. 2018. Recognizing gravity as a strong force in atmosphere emissions markets. *Agricultural and Resource Economics Review* 47 (2): 201–19.

Burtraw, Dallas, Karen L. Palmer, and D. B. Kahn. 2010. A symmetric safety valve. *Energy Policy* 38:4921–32.

Cason, Timothy N., John K. Stranlund, and Frans P. de Vries. 2020. Investment incentives in tradable emissions markets with price floors. Unpublished manuscript, Purdue University, Department of Economics.

Chen, D. L., M. Schonge, and C. Wickens. 2016. oTree: An open-source platform for laboratory, online and field experiments. *Journal of Behavioral and Experimental Finance* 9:88–97.

Dolan, Ed. 2021. Learning from the Green opposition. *Milken Institute Review*, 3rd quarter.

Fankhauser, Samuel, Cameron Hepburn, and Jisung Park. 2010. Combining multiple climate policy instruments: How not to do it. *Climate Change Economics* 1 (3): 209–25.

Fischer, Carolyn, Michael Hübler, and Oliver Schenker. 2021. More birds than stones: A framework for second-best energy and climate policy adjustments. *Journal of Public Economics* 203:104515. https://doi.org/10.1016/j.jpubeco.2021.104515.

Fishelson, Gideon. 1976. Emission control policies under uncertainty. *Journal of Environmental Economics and Management* 3 (3): 189–97.

Flachsland, Christian, Michael Pahle, Dallas Burtraw, Ottmar Edenhofer, Milan Elkerbout, Carolyn Fischer, Oliver Tietjen, and Lars Zetterberg. 2019. How to avoid history repeating itself: The case for an EU Emissions Trading System (EU ETS) price floor revisited. *Climate Policy* 20:133–42. https://doi.org/10.1080/14693062.2019.1682494.

Friesen, Lana, Lata Gangadharan, Peyman Khezr, and Ian A. MacKenzie. 2019. Cost containment in pollution auctions. Discussion Papers Series 610, School of Economics, University of Queensland.

———. 2022. Mind your Ps and Qs! Variable allowance supply in the US regional greenhouse gas initiative. *Journal of Environmental Economics and Management* 112:102620. https://doi.org/10.1016/j.jeem.2022.102620.

Hafstead, Marc, and Roberton C. Williams III. 2017. Adding quantity certainty to a carbon tax through a tax adjustment mechanism for policy pre-commitment. *Harvard Environmental Law Review Forum* 41:41–57.

Holt, Charles A., and William M. Shobe. 2016a. Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the EUETS market stability reserve. *Journal of Environmental Economics and Management* 76:32–50.

———. 2016b. Reprint of: Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the EUETS market stability reserve. *Journal of Environmental Economics and Management* 80:69–86.

Holt, Charles A., William M. Shobe, Dallas Burtraw, Karen L. Palmer, and Jacob K. Goeree. 2007. Auction design for selling CO_2 emission allowances under the Regional Greenhouse Gas Initiative. Final report produced for RGGI, Resources for the Future, Washington, DC.

Isaac, R. Mark, and Charles R. Plott. 1981. Price controls and the behavior of auction markets: An experimental examination. *American Economic Review* 71:448–59.

Karp, Larry, and Christian Traeger. 2018. Prices versus quantities reassessed. University of California, Berkeley, Department of Agricultural and Resource Economics

Khezr, Peyman, and Ian A. MacKenzie. 2018a. An efficient and implementable auction for environmental rights. Discussion Paper Series 587, University of Queensland, School of Economics.

———. 2018b. Permit market auctions with allowance reserves. *International Journal of Industrial Organization* 61:283–306.

Löfgren, Åsa, Dallas Burtraw, Markus Wråke, and Anna Malinovskaya. 2018. Distribution of emissions allowances and the use of auction revenues in the European Union Emissions Trading System. *Review of Environmental Economics and Policy* 12 (2): 284–303. https://doi.org/10.1093/reep/rey012.

MacKenzie, Ian A. 2022. The evolution of pollution auctions. *Review of Environmental Economics and Policy* 1:1–24. https://doi.org/10.1086/717898.

Meckling, Jonas, Nina Kelsey, Eric Biber, and John Zysman. 2015. Winning coalitions for climate policy. *Science* 349 (6253): 1170–71.

Meckling, Jonas, Thomas Sterner, and Gernot Wagner. 2017. Policy sequencing toward decarbonization. *Nature Energy* 2 (12): 918–22. https://doi.org/10.1038/s41560-017-0025-8.

Metcalf, Gilbert E. 2009. Market-based policy options to control U.S. greenhouse gas emissions. *Journal of Economic Perspectives* 12 (2): 5–27.

Montero, Juan-Pablo. 2008. A simple auction mechanism for the optimal allocation of the commons. *American Economic Review* 98 (1): 496–518.

Montgomery, W. David. 1972. Markets in licenses and efficient pollution control programs. *Journal of Economic Theory* 5 (3): 395–418.

Murray, B., R. G. Newell, and W. A. Pizer. 2009. Balancing cost and emissions certainty: An allowance reserve for cap-and-trade. *Review of Environmental Economics and Policy* 3 (1): 84–103.

Newell, Richard, William A. Pizer, and Jiangfeng Zhang. 2005. Managing permit markets to stabilize prices. *Environmental and Resource Economics* 31 (2): 133–57.

Osorio, Sebastian, Oliver Tietjen, Michael Pahle, Robert C. Pietzcker, and Ottmar Edenhofer. 2021. Reviewing the Market Stability Reserve in light of more ambitious EU ETS emission targets. *Energy Policy* 158:112530. https://doi.org/10.1016/j.enpol.2021.112530.

Oswald, A., and N. Stern. 2019. Why are economists letting down the world on climate change? https://voxeu.org/article/why-are-economists-letting-down-world-climate-change.

Pahle, Michael, Dallas Burtraw, Christian Flachsland, Nina Kelsey, Eric Biber, Jonas Meckling, and John Zysman. 2018. Sequencing to ratchet up climate policy stringency. *Nature Climate Change* 8:861.

Perino, Grischa. 2018. New EU ETS Phase 4 rules temporarily puncture waterbed. *Nature Climate Change* 8:260–71. https://doi.org/10.1038/s41558-018-0120-2.

Perino, Grischa, Robert A. Ritz, and Arthur A. van Benthem. 2021. Overlapping climate policies. Working paper 25643, National Bureau of Economic Research, Cambridge, MA. https://doi.org/10.3386/w25643.

Perino, Grischa, and Maximilian Willner. 2019. Rushing the impatient: Allowance reserves and the time profile of low carbon investments. *Environmental and Resource Economics* 74:845–63.

Perkis, David F., Timothy N. Cason, and Wallace E. Tyner. 2016. An experimental investigation of hard and soft price ceilings in emissions permit markets. *Environmental and Resource Economics* 63 (4): 703–18.

Pizer, William A. 2002. Combining price and quantity controls to mitigate global climate change. *Journal of Public Economics* 85 (3): 409–34.

Pizer, William A., and Brian Prest. 2020. Price versus quantities with policy updating. *Journal of the Association of Environmental and Resource Economists* 7 (3): 483–518.

Requate, Till, Eva Camacho-Cuena, Ch'ng Kean Siang, and Israel Waichman. 2019. Tell the truth or not? The Montero mechanism for emissions control at work. *Journal of Environmental Economics and Management* 95:133–52.

Roberts, M. J., and M. Spence. 1976. Effluent charges and licenses under uncertainty. *Journal of Public Economics* 5 (3–4): 193–208.

Rosendahl, Knut Einar. 2019. EU ETS and the waterbed effect. *Nature Climate Change* 9 (10): 734–35.

Salant, Stephen, William Shobe, and Neslihan Uler. 2022. The effects of non-binding price floors. *European Economic Review* 145:104122. https://doi.org/10.1016/j.euroecorev.2022.104122.

Shobe, William, Charles Holt, and Thaddeius Huetteman. 2014. Elements of emission market design: An experimental analysis of California's market for greenhouse gas allowances. *Journal of Economic Behavior and Organization* 107:402–20.

Stern, Nicholas, and Joeseph E. Stiglitz. 2021. The social cost of carbon, risk, distribution, market failures: An alternative approach. National Bureau of Economic Research, Cambridge, MA.

Stranlund, John K., James J. Murphy, and John M. Spraggon. 2014. Price controls and banking in emissions trading: An experimental evaluation. *Journal of Environmental Economics and Management* 68 (1): 71–86.

US Department of State and US Executive Office of the President. 2021. The long-term strategy of the United States: Pathways to net-zero greenhouse gas emissions by 2050.

Weitzman, Martin L. 1974. Prices vs. quantities. *Review of Economic Studies* 41 (4): 477–91.

Recognizing Gravity as a Strong Force in Atmosphere Emissions Markets

Dallas Burtraw and Amelia Keyes

Environmental economics has made it possible to estimate prices for air pollution externalities. However, these values are rarely observed in emissions trading markets. Moreover, market outcomes show prices persistently remain below expectations and frequently fall over time. Low allowance prices may appear virtuous, but often reflect poor market design that does not anticipate interaction with other policies, and may undermine confidence in market-based approaches to environmental policy. This paper surveys emissions markets and factors influencing prices, and concludes with a discussion of how market design can anticipate and remedy the strong tendency for low prices.

Key Words: allowance prices, cap-and-trade, cost management, innovation

JEL Classification Numbers: D47, Q02, Q58

A major contribution of environmental economics has been the technical ability to estimate a price for the marginal value of nonmarket goods and services, and in particular, to value the injuries to environmental services that result from the emission of pollutants. The generally recommended approach for bringing this information to policy is the damage function approach, an integrated assessment that, for atmosphere resources, links emissions to atmospheric transport of pollutants to pollutant exposure to injury from exposure and ultimately to valuation of the injury. The estimate of total damages represents the change from a reference case representation of the service flows in absence of the emissions. This quantification of damages provides an evidentiary basis for applying the polluter pays principle—in which the emitter is required to bear the costs of the pollution—to atmosphere emissions based on the change caused by emissions in the value of environmental services. This information should inform the quantity of emissions allowed under regulations such as cap-and-trade, or the emissions price that might be charged under a fee or tax approach.

However, lest we allow the perfect to be the enemy of the good, this framework for public policy has not been used directly in environmental regulation. Even in cases where market-based approaches have been used,

---

Dallas Burtraw is a senior fellow and Amelia Keyes is a research assistant at Resources for the Future. *Correspondence: Dallas Burtraw ▪ Resources for the Future ▪ 1616 P St. NW ▪ Washington, DC 20036 ▪ Email: burtraw@rff.org.*

The views expressed are the authors' and do not necessarily represent the policies or views of any sponsoring agencies.

damage function estimates have not determined regulatory stringency. Instead, this scientific information has simply been a data input to regulatory negotiations that balance a multitude of other political considerations.

There are two dominant observations to make about the experience with market-based approaches to atmosphere resource regulations. First, emissions targets adopted in these environmental markets have usually been higher (less stringent) than would be recommended by a straightforward balancing of benefits and costs. Among many factors contributing to this outcome, it is sufficient to note the logic of collective action, which prescribes that diffuse benefits from environmental improvement will be devalued in the political context compared to the concentrated costs that accrue to well-identifiable and consequently well-motivated opponents (Olson 1965). This appears to apply especially well when environmental policy instantiates economic value in a new asset by introducing a scarcity of tradable emissions allowances. The typical experience is that the regulatory goal in atmosphere markets is initially set at an emissions target where the marginal benefits of additional reductions would be greater than the marginal costs, i.e., the regulation is too lax from an economic viewpoint. Furthermore, new information in science and economics has tended to cause the estimated benefits of emissions reductions to increase and the balance of benefits and costs to evolve even further in the direction of low stringency. Even where there has been approximate balancing of marginal benefits and marginal costs in setting the initial environmental target, for example, as may have applied in the case of the 1990 Clean Air Act emissions cap for sulfur dioxide (Portney 1990), the revelation of new scientific and economic information typically has led the initial assessment of the benefits of emissions reductions to rapidly become out of date and be revised upwards (Chestnut and Mills 2005).

The second observation, which is the focus of this essay, concerns the costs of achieving the stated emissions target. In atmosphere resource markets, the marginal cost of achieving the target is manifest in emissions allowance prices. Historically, emissions allowance prices consistently have fallen below the values anticipated ex ante when the regulations were initially promulgated. We describe this factor as a *strong force* in the title of this essay to indicate, by way of a mash-up analogy to the fundamental interactions of physics, that the downward tendency for prices in allowance markets persists strongly even in the face of other factors that affect market dynamics.

Lower-than-expected prices may appear less like a problem than a virtue because they may represent lower-than-expected costs of compliance. However, programs that do not behave as expected by policymakers or as described in economic theory can pose a political problem. Moreover, if programs behave in a way not described by theory then it undermines the credibility of economic perspectives. As science and policy lead us toward deep decarbonization of the economy, the design of environmental markets must account for the likelihood of declining prices if economic tools and their associated advantages will be influential in achieving environmental goals.

The ability of markets to anticipate and respond to this tendency in allowance prices has proven essential to their durability. The primary response of successful programs involves enforcing a minimum price, implemented using a reserve price in emissions allowance auctions, as occurs in all the North American carbon dioxide (CO_2) emissions markets. More generally, a strong program design would include a price-responsive (elastic) schedule for introducing emissions allowances into the market; that is, an upward sloping supply schedule. All North American programs provide for *increases* in supply at specified (high) trigger prices, but except for the provision of a minimum price, providing for decreases in supply at low trigger points is less prevalent.

Like the other North American programs, the Regional Greenhouse Gas Initiative—a carbon trading program in nine northeast US states—has adopted both a minimum price and a schedule for introducing allowances into the market at high prices. In addition, it adopted a third feature, a schedule to *decrease* the supply of allowances in the market at low allowance prices, an approach it describes as an *emissions containment reserve*. After a decade of prices that trended well below expected levels, the EU's CO_2 Emissions Trading System adopted another approach to respond to the downward tendency in allowance markets by back-loading (delaying) auctions and prescribing in 2019 the movement of allowances into a reserve in response to allowance surpluses. After that approach had little effect on current allowance prices, in 2018 the EU introduced another provision for cancelling allowances beginning in 2023 if the surplus grows large, which has in fact triggered a response with an increase in prices.

This essay surveys the performance of emissions allowance trading programs and the trend for emissions allowance prices to fall below initial expectations. We describe several factors influencing that trend and conclude with a discussion of how market design can take this strong tendency into account.

The Political Narrative Has Focused on Potentially High Prices

Our premise is that the dominant feature of emissions allowance markets has been a downward trend in allowance prices. However, ex ante, the possibility of a price increase has dominated the political narrative when introducing emissions markets. Like Rougarou, the seldom-seen mythical beast of the French alps, a price spike has been observed rarely, and a sustained price increase has never been witnessed in emissions allowance markets – at least not yet. Nonetheless, if one were representing the interests of a firm with a compliance obligation and exposure to a potential price spike that might dramatically increase costs, one would advocate for a program design to mitigate that risk.

There is good reason for prices to be expected to rise. A fundamental observation from economic theory is that an efficient policy will achieve the least-cost emissions reductions first, a result facilitated by allowance trading.

If it were less expensive to reduce emissions at a facility than to use an allowance valued at the market price to cover incremental emissions, the firm should reduce its emissions and sell the saved allowance back into the market (or not purchase an additional allowance); if it were more expensive, then the firm would save money if it were to buy an allowance from the market rather than reduce its own emissions. This process leads the marginal cost across all firms to equilibrate and equal the market price of an allowance. As the stringency of an emissions cap is increased (the emissions cap is lowered), the marginal cost of achieving emissions reductions should increase also. All the existing atmosphere resource programs exhibit increasing stringency (lower caps) over time, and so prices might be expected to rise in those programs.[1] Moreover, firms might be especially concerned about unanticipated price increases that might result, for example, from disruptions in related energy factor markets. For example, the experience in the sulfur dioxide market shows that Hurricane Katrina, which knocked out natural gas supply and boosted coal use, led to a short-term spike in allowance prices.

Several features in emissions markets help mitigate the chance of sudden price spikes. One of the most important is emissions allowance banking, which enables an allowance to be eligible for use in compliance at any time after its initial issuance so that firms can ration their compliance activities over time. A firm might choose to make a major investment in emissions control and over-comply with its compliance obligation, knowing that an allowance saved is a dollar earned – that is, allowances not used for compliance in the current compliance period will continue to hold their value for future compliance periods and can be expected to increase in value if compliance costs across the industry increase. The ability to bank conveys convenience value to the firm, not only in scheduling investments and production activities, but also hedging against a potential spike in allowance prices. Banking converts a single-year compliance obligation into a multi-year compliance horizon for planning purposes, yielding an expanded pool of allowances to use for compliance in the face of any short-run disruptions in relevant factor markets. In fact, where a short-term runup of allowance prices has occurred, it has occurred in programs that lacked banking.

Economists have often advocated for the opportunity to borrow allowances from a future compliance period as another way to mitigate a potential price spike, but the challenges of moral hazards are severe. For instance, the least solvent firm in an industry might be most likely to borrow on its future allocation of allowances and sell them into the market, and if the firm ever

[1] If there is a positive number of allowances in privately held banks, this also suggests the allowance price increase over time, even if the initial introduction of emissions allowances (the cap) is constant over time. The allowance price will rise, reflecting the opportunity cost of capital invested in holding that allowance, compared to the option of selling it into the market.

becomes insolvent the obligation might not be repaid. However, most programs enable implicit short-term borrowing through multi-year or overlapping compliance periods, another method for softening the impact of short-term pressures on the allowance market.

The carbon markets also allow for the use of out-of-market compliance, or *offsets*, which provide an additional way to achieve compliance. By bringing additional compliance instruments into the program, the introduction of offsets effectively expands the supply of compliance instruments and lowers allowance prices in the market.

Where all else fails, the existing carbon programs have a reserve of allowances that enter the market and increase the supply of allowances at a trigger price. California's extension to its cap-and-trade program through 2030, which takes effect in 2021, introduces for the first time a hard price ceiling at a level to be determined by the state's Air Resources Board, at which an unlimited supply of allowances would become available in the market. This event would not supersede the Board's obligation to achieve statutory requirements to achieve specified emission reductions, but they would have to occur outside the market. California and the RGGI program also incorporate a cost-containment reserve, a quantity of allowances that are made available if allowance prices exceed a predetermined level.

In summary, there are a number of reasons why allowance prices typically are expected to rise over time in most programs, and this expectation is built into the policy design. Moreover, policymakers and compliance entities often focus their concern on the possibility of allowance prices rising at an even faster rate than anticipated. As described above, this political narrative has had a substantial impact on program design: allowance markets generally integrate a set of mechanisms to protect against the risk of price spikes. However, the prevalent trend is not rising prices but lower-than-expected or even falling prices. The following sections will highlight some of the reasons for this outcome and describe mechanisms that can be used in allowance markets to address this reality.

The Actual Outcome in Markets

The experience across atmosphere resource markets has not been rising prices, and only rarely have the markets experienced short-run price spikes. Across all these markets, the trend has consistently been lower than expected and frequently falling prices.

US Sulfur Dioxide Market

The first grand experiment in emissions allowance trading was the trading of sulfur dioxide allowances in the US electricity industry, authorized by the 1990 Clean Air Act Amendments. Sulfur dioxide (SO_2) contributes to acidification of terrestrial ecosystems and to formation of fine particulates, which affect public health. The first phase of the program began in 1995 and affected 110 of the largest and most polluting coal-fired generating stations. In the second phase, beginning in 2000, the program expanded to cover the entire industry. Banking of allowances over time and between phases was allowed. Virtually all allowances were distributed for free, based on historic use of facilities, but 2.8 percent of the allowances were withheld and distributed through a consignment auction, with revenues distributed proportionately back to the facilities. The first bilateral trades occurred at prices between \$200 and \$400 per allowance (Ellerman et al. 2000). The first annual auction, held in 1993, helped identify a market price and stimulated transactions. Over the first six years, auctions cleared at prices ranging from \$70 to \$190, and bilateral trades in the market fell in the same neighborhood, well below the band of values forecasted by the US EPA and market analysis, shown in the yellow band in Figure 1 (ICF 1990). The figure also displays the draft and final versions of the Clean Air Interstate Rule (CAIR) that fundamentally restructured the market by changing the value of allowances that would be issued after 2010, and the price spike associated with Hurricane Katrina, which disrupted natural gas facilities and led to an increase in coal use. From the first year of compliance in 1995 until the introduction of CAIR, marginal costs and allowance prices were one-third to one-quarter the values anticipated before implementation, and prices fell again after CAIR provisions pertaining to SO_2 were suspended by the courts. A multitude of reasons contributed to the lower-than-expected prices, including the expanded availability and low cost of low-sulfur coal, falling prices for post-combustion controls (scrubbers), and unexpected success in blending of low- and high-sulfur fuels (Burtraw and Szambelan 2009). These innovations played an important role, but the low prices were also due in a major way to the dynamic flexibility of the program design, which allowed

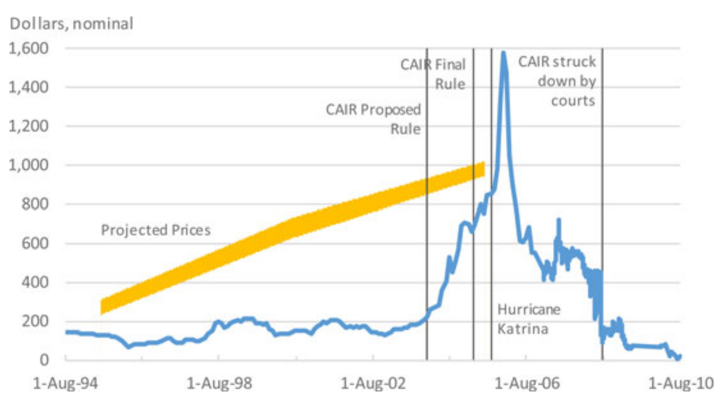

Figure 1. US sulfur dioxide allowance prices

https://doi.org/10.1017/age.2018.12 Published online by Cambridge University Press

compliance entities to use a broad range of abatement options and to bank their allowances, thus enabling the allowance market to capture the benefits of trends in the electricity and related factor markets (Burtraw and Szambelan).

US Nitrogen Oxides Market

The second important trading program regulated emissions of nitrogen oxides from the power sector in 19 states in the eastern United States. Nitrogen oxides (NO_x) are a precursor to ozone formation, a summer season problem, and contribute to the formation of fine particulates and acidification on a year-round basis. This program was applied during the five-month ozone season and was the outgrowth of a previous program involving only northeastern states. It was precipitated by the US Environmental Protection Agency's call for states to revise their state implementation plans, as shown in Figure 2. The late entry of Maryland to the trading program and a dramatic increase in natural gas prices caused a short-term price spike at the outset of the program, but prices soon fell to a fairly stable level with the widespread adoption of selective catalytic reduction, and boiler-level adjustments in facility operations (Burtraw and Szambelan 2009). The advent of the CAIR program for nitrogen oxides drove further investments in emissions reductions and further depressed allowance prices.

Chicago ERMS for Volatile Organic Compounds

A third example is the local Emissions Reduction Market System (ERMS) in the Chicago area for volatile organic compounds (VOCs), precursors to formation of ozone (Figure 3). This program began in 2000 and was designed to be a lower-

Figure 2. Regional nitrogen oxide emission budget prices

cost approach to reduce VOCs emissions in the Chicago ozone nonattainment area. Similar to the NO_x program, ERMS compliance applies only during the summer ozone season. Contrary to expectations, emissions in the first years of the program were substantially lower than their annual caps; consequently, allowance prices declined and were far lower than originally predicted—in 1996 the Illinois EPA predicted allowance prices of $344, and since 2003 the average annual allowance price has been at or below $20 (Evans and Kruger 2007). Due to the lack of allowance scarcity, minimal trading occurred, and many allowances expired without being used. Unused allowances and lower-than-expected allowance prices were attributed in part to an initial over-allocation of allowances that failed to create scarcity. Allowances were distributed at no cost to compliance entities based on their highest summer emissions, and the cap was set as the sum of allocated allowances. This method is distinct from the common method in which an emissions cap is predetermined and allowances are distributed based on sources' historical share of total emissions. There is evidence suggesting that baseline emissions in ERMS were inflated; emissions in 1998 before the program began were only 67 percent of the total allocated emissions in 2000, the first year of the program (Evans and Kruger). Throughout the life of the program, emissions may also have been driven down by other regulations, including Illinois state regulations for ozone compliance and federal controls for hazardous air pollutants (Kosobud et al. 2004). Compliant sources also experienced a number of shutdowns over the course of the program and continued to receive allowance allocations, putting further downward pressure on allowance prices. In 2005, 36 sources with no emissions were still receiving allowance allocations (Evans and Kruger).

Figure 3. Local volatile organic compound prices

California RECLAIM for Sulfur Oxides and Nitrogen Oxides

Another example of a local emissions trading program is the Regional Clean Air Incentives Program (RECLAIM) in southern California, which is the exception that establishes the rule. It is the one program that had a spectacular increase in prices that led to the temporary suspension of the program. The RECLAIM program began in 1994 and created markets for both sulfur oxides and NO_x emissions. For the first six years of the program, emissions by compliance entities were below the declining annual cap, and allowance prices remained low and stable (US EPA 2002). However, in 2000–2001 the prices for NO_x allowances in RECLAIM soared, interacting with the dry hydro year and the California electricity crisis. NO_x allowance prices in compliance year 2000 reached above $45,000 per ton, over ten times the average price in 1999 (US EPA). The key feature of the RECLAIM program that enabled this outcome was the absence of banking. Compliance entities were organized in two groups with one-year compliance periods that overlapped by six months. Consequently, there was a limited ability to purchase allowances from a compliance entity with a different compliance period and effectively borrow from the short-term future. In practice, the absence of banking led to "just-in-time" emissions reductions. That design, coupled with increasing stringency of the program, a low hydro year, other changes in the electricity markets, and intentional market manipulation, created the circumstances that led to the program's collapse. The experience shows that, just as a poorly designed house will fall down, a poorly designed program might collapse also.

EU ETS for Carbon Dioxide

The first trading program for CO_2 took shape in the EU's emissions trading system (ETS), which began in what has been called a trial period from 2005–2007. The program covers about 11,000 stationary facilities accounting for about 45 percent of EU emissions. The EU ETS Phase 1 had a design feature in common with RECLAIM—it broadly limited banking of allowances into Phase 2 of the program. In a program without banking, towards the end of the compliance period when allowances become worthless two things can happen. Prices may go to zero because allowances the day after the end of the compliance period are worthless, as was observed in the EU ETS (Figure 4). Or allowance prices may go toward infinity because supply is exhausted, as in RECLAIM. The EU ETS and RECLAIM illustrate these two possibilities coming to life.

In the subsequent Phases 2 (2008–2012) and 3 (2013–2020) of the EU ETS, allowance banking has been allowed. The relaunch of the program in Phase 2 led to an initial recovery of prices, approaching the level forecasted by analysts as necessary to achieve the EU's long-term climate goals. However, those prices have fallen substantially and consistently over the program's history. Several factors contributed to this price trend, including the

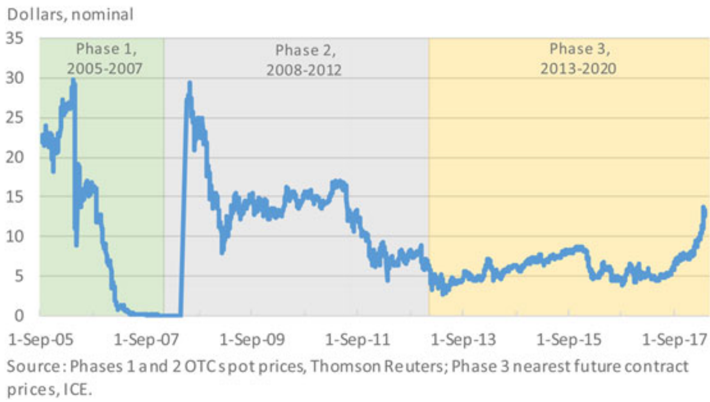

Figure 4. Prices in the EU carbon dioxide emissions trading system

recession, the availability of offsets through the international Clean Development Mechanism, and the substantial development of renewable energy promoted by government programs (Knopf et al. 2014, Bel and Joseph 2015). These factors have driven down the demand for emissions allowances, and allowance prices have fallen accordingly. With the ability to bank allowances beginning in Phase 2, allowance prices have reflected not just the relationship between allowances and emissions in the current period, but also the expectations of future allowance scarcity (Knopf et al. 2014). Persistently low allowance prices motivated repeated reforms to the program, and beginning in 2018 the most important of these reforms had its desired effect, leading to an important uptick in prices.

Northeastern US RGGI Carbon Dioxide Market

RECLAIM notwithstanding, because it did not allow banking, the general trend across programs discussed so far was falling prices. This observation was evident and supported two program design innovations in the Regional Greenhouse Gas Initiative (RGGI) CO_2 trading program. One innovation was a central role for auctions in the initial distribution of emission allowances, motivated by observations in the EU ETS of the transfer of wealth and windfall profits to emitters due to free allocation. The second innovation was the introduction of a minimum price, enforced through what is called a reserve price in the quarterly auctions. A reserve price prohibits allowances from being sold at any price below an established minimum price. This works just like the option to specify a minimum price on eBay; it constricts the supply of allowances when prices fall to the identified reserve price, thereby supporting the price in the market.

In RGGI, the auction reserve price has proven crucial to the survival of the program. After the first few auctions, the auction price settled at the reserve price for eleven consecutive auctions before changes in the program led to a recovery of prices. The reserve price was low, at about $2 per ton, illustrated as a "price floor" in Figure 5. However, the reserve price helped the auction net over a billion dollars in revenue in that period, which were primarily invested in program-related activities such as energy efficiency. By maintaining an active market and preserving the value of early investments and banked allowances, the reserve price provided buoyancy to the program until scarcity was adjusted to lift prices off the minimum price level in the auction.

The experience with the RGGI program offered a model for the proposed national trading program known as the Waxman-Markey proposal, which passed the US House of Representatives in 2009. That proposal included an important and growing role for an auction, with a reserve price in the auction beginning at $10 per ton and increasing over time. That proposal was never adopted into law, but the model was propagated further in the California program, illustrated in Figure 6.

California Carbon Dioxide Market

The first trades in the California CO_2 market occurred in 2011, auctions began in 2012, and compliance began in 2013, covering only the electricity sector and large stationary sources. In 2015 it expanded to an economy-wide program, covering about 85 percent of emissions in the state, linked with Quebec, and subsequently linked with Ontario in 2018. The program has a reserve price

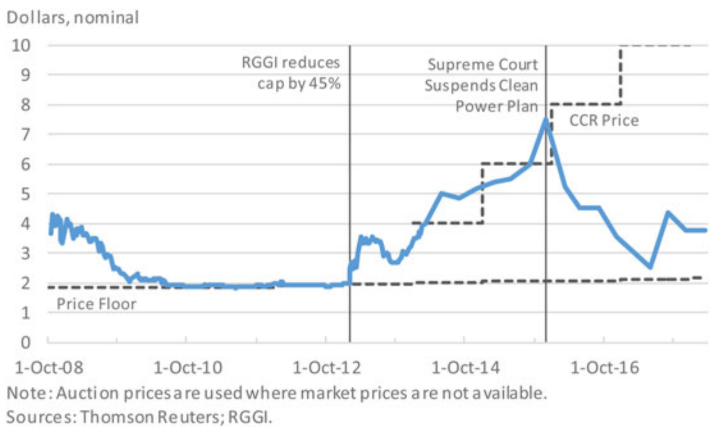

Figure 5. Prices in the Regional Greenhouse Gas Initiative carbon dioxide program

https://doi.org/10.1017/age.2018.12 Published online by Cambridge University Press

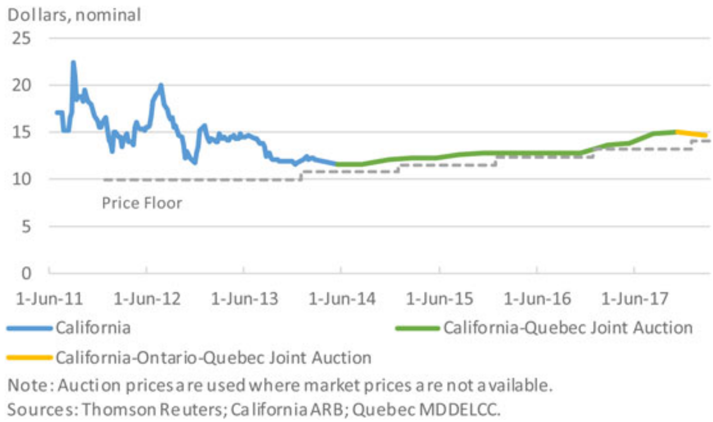

Figure 6. Prices in California, Quebec and Ontario carbon dioxide cap-and-trade programs

in its auctions that increases at 5 percent per year plus inflation. The reserve price caused some allowances to remain unsold in five consecutive auctions in 2016, before demand stiffened and the auction supply was again fully sold.

Factors that Explain Downward Trends in Allowance Prices

We argue that the most important and consistent phenomenon in emissions trading programs has been falling allowance prices. The next question follows – why? A number of factors contribute to the "gravitational" force witnessed so often in trading programs.

The first and most obvious cause of low prices is over-allocation of allowances. The determination of annual emissions caps and allowance allocation is subject to political influence. Frequently, emissions caps are set at levels that are less stringent than would be suggested if emissions caps were informed by the damage function approach to measuring marginal benefits of emissions reductions. This was particularly evident in the Chicago ERMS program, where the emissions cap was based on the sum of compliance entities' baseline emissions. Evidence suggests that baseline emissions in ERMS were inflated, leading to an emissions cap that required minimal abatement and enabled low allowance prices (Evans and Kruger 2007).

In some cases, low allowance prices reflect program uncertainty. If compliance entities do not have confidence in a program's future, they may bet against it and reduce their demand for allowances. Knopf et al. (2014) found evidence suggesting that political factors and regulatory uncertainty played the most important roles in the price declines in earlier phases of the EU ETS.

Allowance prices often fall below expectations because the incentives created by emissions trading programs work as intended. Burtraw and Szambelan (2009), in their analysis of the SO_2 trading program, found that the costs of abatement fell over time in part because trading triggered a search for ways to reduce emissions at lower costs. The trading program provided flexibility in the abatement methods that compliance entities could use, such as fuel blending and switching, which allowed them to fully exploit innovations and trends in fuel and technology markets. Additionally, flexibility allowed for competition among abatement options, which helped to further reduce compliance costs (Burtraw 1996, 2000). Analysis of the NO_x Budget Program found similar results. Under this program, an estimated 10 to 15 percent of emissions reductions resulted from small-scale modifications to existing capital that would not likely have occurred under a command-and-control policy (Linn 2008, Burtraw and Szambelan 2009). These innovations, along with fuel switching and the retirement of noncontrolled units, contributed to lower compliance costs and falling allowance prices. In general, innovations such as these are difficult to anticipate, so they often cannot be enumerated in advance, and hence they are often not reflected in modeling that provides ex ante estimates of the cost of compliance and that influences decisions about the stringency of the emissions caps.

Another driver of falling allowance prices is subjurisdictional policy. Jurisdictions within a trading regime may choose to implement more stringent policy measures, which can reduce the total demand for allowances. For instance, in the EU ETS several member states, including Sweden, Norway, Denmark, the UK, and Ireland, have implemented domestic carbon taxes or other forms of carbon pricing on top of the international emissions trading scheme. Many member states also provide significant subsidies to renewable energy that affect allowance prices (Van den Bergh, Delarue, and D'haeseleer, 2013, Koch et al. 2015).

Finally, and importantly, emissions trading programs exist within a landscape of other policies and companion policies, such as regulatory standards and government investment, that have a significant effect on the outcomes of trading programs. These companion policies are more often called "complementary policies" by economists who envision emissions pricing as the primary policy. However, in every case, the pricing policies follow earlier regulatory efforts that employ standards or promote technologies, and these policies are typically maintained and sometimes strengthened with the advent of emissions pricing (Meckling, Sterner, and Wagner 2017, Pahle et al. 2017). Hence, we might ask whether or not emissions pricing is a complementary policy. In any event, we choose to call them *companion policies* because they always are apparent wherever emissions pricing is observed.

Companion policies often exist to address concerns beyond the targeted type of emissions. These policies can address air quality, job creation, economic development, institutional development, and good old-fashioned fighting for rents. Companion policies can be particularly useful for jurisdictions that take leadership positions on addressing emissions and are concerned about leakage, because companion policies can have a smaller effect on the changes in final prices of goods and services, which can help preserve the competitiveness of trade-exposed industries. However, companion policies always reduce the demand for emissions allowances and push down prices. In the EU ETS, the use of offsets through the Clean Development Mechanism and the development of renewable energy encouraged by government programs have put downward pressure on allowance demand and prices. In California, companion policies played an early role and continue to play a major role in meeting emissions reduction goals and reduce the cost of compliance in the cap-and-trade program. The state's first and second emissions reduction scoping plans expected that the regulatory standards and measures were sufficient to achieve over 80 percent of their targeted emissions reductions, with the cap-and-trade program accounting for the remainder of reductions. Their recent third scoping plan expects regulatory standards and measures to continue playing a major role by achieving 60 percent of reductions, and leaving the remaining 40 percent to cap and trade. Similarly, complementary policies are a prominent feature of the RGGI program, where over 85 percent of the revenue from the auction of allowances has been invested in program-related spending including energy efficiency programs and clean energy investment. These program-related investments are designed to reduce the demand for allowances and keep allowance prices low. The RGGI program is given credit for about half of the total emissions reductions observed in the region over its timeframe (Murray and Maniloff 2015).

Trading programs have also been influenced by exogenous conditions such as technology changes and economic shocks. The greatest of these exogenous influences on CO_2 emissions markets in recent years were the economic recession, and in the United States, the shale gas boom. The economic recession reduced demand for goods and services and therefore decreased baseline emissions, putting downward pressure on allowance prices. The shale gas revolution led to a greater supply of low-cost natural gas, which, because of its low emissions profile compared to coal, provided a lower cost emissions abatement option and put downward pressure on allowance prices. Although these phenomena contributed to declining prices, it is equally possible for exogenous economic and technology changes to have an increasing effect on allowance prices – but this is not what we have observed.

What are the Consequences of Falling Prices?

As we discussed above, the possibility of rising allowance prices is one that has been discussed extensively, and mechanisms are designed into emissions trading programs to address it. However, the trend of falling prices is far more prevalent and merits a discussion of the consequences and the mechanisms that can help avoid it.

One might ask, what is wrong with a low price if the emissions target is being achieved? The concern is that low allowance prices may fail to create price incentives that facilitate investments that lead to reductions in an economy's emissions intensity over the long term. Low prices, for example, may not trigger changes such as fuel switching or renewables uptake, because compliance entities can purchase allowances at a lower price than would be required to make these investments. This is an issue even if the emissions trading program is meeting its short-term emissions targets. For example, Knopf et al. (2014) found that the allowance price path in the EU ETS has not matched a socially optimal CO_2 price path that would achieve climate targets at the least cost over time.

There is another implication of prices that fall well below expectations – companion policies may not have their anticipated effect on the emissions outcome. Companion policies, such as subjurisdictional policy efforts within regional programs in the EU or RGGI, do not change the emissions cap. Instead, they change the profile of emissions reductions that would occur if the emissions price were the sole influence on behavior, and under a cap they lead to 100 percent leakage among jurisdictions (Goulder and Stavins 2011). This is known as the *waterbed effect*, which describes the fact that when you push down on emissions in one place it leads to an increase in emissions in another place, without any change in total emissions. This outcome is insidious because, as we have argued, companion policies have been essential in developing the opportunity for emissions pricing to emerge initially and continue to play a role in price emission jurisdictions, especially carbon. At the same time, however, the waterbed effect may have the perverse effect of undermining the effectiveness of other policies that remain necessary to achieve climate policy outcomes.

A New Generation of Market Designs?

Falling prices in previous emissions markets, especially in the EU ETS, led the architects of the North American carbon markets to innovate in designing a minimum auction price into the market, which we have described previously. Although the minimum price was binding in 11 consecutive auctions, RGGI was nonetheless also confronted with accumulation of a large, privately held allowance bank. In its first program review, RGGI adjusted the quantity in the market by cancelling unsold allowances and by reducing the number of new allowances entering the market to absorb the surplus of allowances then in circulation.

In 2017, RGGI issued an updated model rule to take effect in 2021 that adds an important new program element: the emissions containment reserve (ECR). The ECR introduces a price step below the expected market price and above the price floor. A quantity of allowances (10 percent of the annual cap) would not be sold for a price below the trigger price on the price step ($6 per ton, rising at 7 percent per year). If these allowances are not sold, the price could fall below $6, which is why we refer to the ECR as a soft price step. The unsold allowances are permanently canceled. The price could still fall to the price floor, which applies to all other allowances and provides an absolute minimum price in the auction. The RGGI design is illustrated in Figure 7, where a price below the ECR leads to a reduction in supply. Note also the cost-containment reserve, that in parallel fashion would introduce a quantity of additional allowances at a trigger price above the expected price. In principle, additional price steps could be added to make the allowance supply schedule increasingly responsive to the market price. A price-responsive supply schedule is a characteristic of commodity markets; for example, if the price of natural gas declines, less natural gas enters the market. This characteristic helps reduce price volatility (Burtraw et al. 2018). In emissions markets, however, supply has usually been inelastic and has not varied with respect to the allowance price.

What is the consequence of the falling demand for emissions allowances under the ECR? It is manifest in a lower-than-expected price, which works to the benefit of compliance entities and perhaps economic interests more broadly. But, there also is a contraction in the supply of allowances that limits the fall in price and works to the benefit of the environment by reducing emissions below the emissions target. Consequently, there is a "sharing" of the advantageous gains resulting from falling prices, which is different from the way emissions markets have worked previously. The result is a hybrid instrument, combining elements of an emissions tax with cap-and-trade, an approach that has long been suggested in the economic literature (Roberts and Spence 1976). A program might consider one price step, several

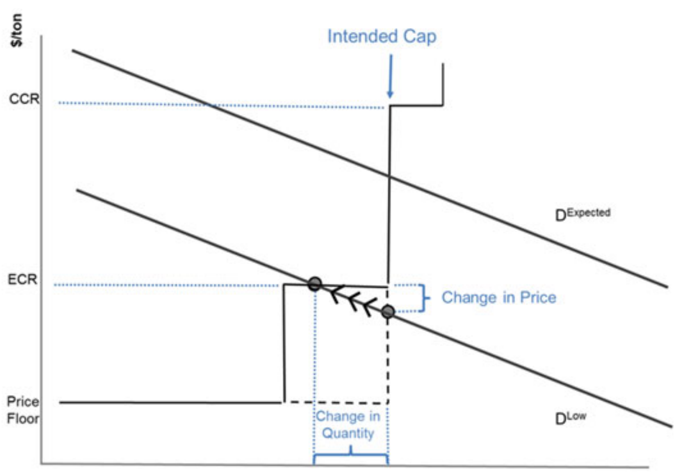

Figure 7. The Emissions Containment Reserve in the RGGI market

price steps, or if a modeler will see right away that the best design (because it is easiest to model!) is a ramp, which would provide a continuous adjustment to both quantity and price.

The EU ETS has recently adopted a different, quantity-based approach, to address their challenge of a low allowance price accompanied by a large emissions bank. In 2015, the EU delayed the issuance of a large volume of allowances and introduced the Market Stability Reserve (MSR), which would automatically restrict the quantity of allowances to be sold in the auction and direct the unsold allowance into the MSR when the quantity of allowances in circulation exceeds a threshold value. Allowances held in the MSR could be returned to circulation if the number of allowances fall below a low threshold value. Analysts were generally skeptical of the effectiveness in delaying the sale of allowances. However, because with banking, the allowance price in the present should anticipate the expanded supply that would ultimately occur (Hepburn et al. 2016). In 2018, the EU took an important further step in prescribing the cancellation of a portion of allowances from the MSR each year when the MSR exceeds a specified volume. This approach appears to have immediately affected the allowance price in the EU (Figure 4), and analysis suggests it will partially, but not fully, remedy the waterbed effect (Perino 2018).

Changes may also follow in California, where the allowance price is supported by the minimum auction price, but nonetheless a substantial quantity of allowances have accumulated in public and private accounts. In legislative testimony in January 2018, the chairperson of the California Air Resources Board acknowledged the need to make sure that the trading program is sending the message to promote investments that are going to reduce emissions over the long run.[2] As we noted above, cap-and-trade is expected to be twice as influential relative to other regulatory measures in driving emissions reductions over the next decade in California. One might interpret this to suggest that regulators will consider measures to enable allowance prices to increase.

Conclusion

In review, we argue that a strong force in emissions allowance markets is gravity – that is, a tendency for prices to fall over time. There are several reasons for this, and they are not all bad. Furthermore, emissions markets that have anticipated falling prices in their program design have been successful; those that have not have been less successful.

It is critical to note that our observation of gravity as a strong force does not imply that prices will always fall. Perhaps with new levels of ambition for

[2] http://calchannel.granicus.com/MediaPlayer.php?view_id=7&clip_id=5086 at 1:09:50.

emissions abatement, prices will rise, even above expectations. Indeed, the programs with a price floor also have a cost-containment mechanism to guard against price increases. These cost-containment reserves take the form of price steps, implemented symmetrically to the idea of an emissions containment reserve that we have described. At specified trigger prices, additional allowances enter the program.

Combining the cost-containment reserve for prices above expected levels with an emissions containment reserve for prices below expected levels maps out a price-responsive allowance supply schedule that might better reflect the politically negotiated scarcity that we observe in market-based programs. We think of the emissions cap as an instruction to the market, but if it turns out that emissions reductions are less expensive than anticipated, then regulators might be expected to purchase more of them. Emissions markets are an important effort to try to take advantage of the power of incentives to achieve environmental outcomes in an efficient way. One might observe that program adjustments are a proven way to adjust supply in response to unanticipated market conditions, but relying on administrative interventions introduces an unpredictable element that accentuates uncertainty associated with environmental markets. Design elements such as an auction reserve price and emissions containment reserve, and potentially the market stability reserve, are rule-based mechanisms that are built into trading programs and automatically adjust the market to deal with low allowance prices. We view these developments as the most encouraging innovation in emissions markets since the widespread introduction of auctions, in lieu of free allocation a decade ago, and the improved design significantly expands the potential for environmental markets.

References

Bel, G., and S. Joseph. 2015. "Emission Abatement: Untangling the Impacts of the EU ETS and the Economic Crisis." *Energy Economics* 49(C): 531–539.

Burtraw, D. 1996. "The SO_2 Trading Program: Cost Savings without Allowance Trades." *Contemporary Economic Policy* 14: 79–94.

——— 2000. *Innovation under the Tradable Sulfur Dioxide Emissions Permits Program in the U.S.* Electricity Sector. Discussion paper 00-38. Washington, DC: Resources for the Future.

Burtraw, D., C. Holt, K. Palmer, A. Paul, and W. Shobe. 2018. "Quantities with Prices." RFF Working Paper 18-08.

Burtraw, D., and Szambelan, S.J. 2009. *U.S. Emissions Trading Markets for SO_2 and NO_x*. Discussion Paper 09-40. Washington, DC: Resources for the Future.

Chestnut, L.G., and Mills, D.M. 2005. "A Fresh look at the Benefits and Costs of the US Acid Rain Program." *Journal of Environmental Management* 77(3): 252–266.

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

Evans, D.A., and Kruger, J.A. 2007. "Where are the Sky's Limits? Lessons From Chicago's Cap-and-Trade Program." *Environment: Science and Policy for Sustainable Development* 49(2): 18–32.

Goulder, L., and R. Stavins. 2011. "Challenges from State-Federal Interactions in US Climate Change Policy." *American Economic Review: Papers & Proceedings* 101(3): 253–257.

Hepburn, C., K. Neuhoff, W. Ackworth, D. Burtraw, and F. Jotzo. 2016. "The Economics of the EU ETS Market Stability Reserve." *Journal of Environmental Economics and Management* 80: 1–5.

ICF 1990. "Comparison of the Economic Impacts of the Acid Rain Provisions of the Senate Bill (S. 1630) and the House Bill (S. 1630)." Prepared for U.S. Environmental Protection Agency, Washington, DC (July).

Knopf, B., N. Koch, G. Grosjean, S. Fuss, C. Flachsland, M. Pahle, M. Jakob, and O. Edenhofer. 2014. "The European Emissions Trading System (EU ETS): Ex-Post Analysis, the Market Stability Reserve and Options for a Comprehensive Reform." Working Paper 79.2014. Milan, Italy: Fondazione Eni Enrico Mattei.

Koch, N., S. Fuss, G. Grosjean, and O. Edenhofer. 2015. "Causes of the EU ETS Price Drop: Recession, CDM, Renewable Policies or a Bit of Everything?—New Evidence." *Energy Policy*, 73: 676–685.

Kosobud, R.F., H.H. Stokes, C.D. Tallarico, and B.L. Scott. 2004. "The Chicago VOC Trading System: The Consequences of Market Design for Performance." MIT CEEPR Working Paper WP-2004-019. Boston.

Linn, J. 2008. "Technological Modifications in the Nitrogen Oxides Tradable Permit Program." *The Energy Journal* 29(3): 153–176.

Meckling, J., T. Sterner, and G. Wagner. 2017. "Policy Sequencing Toward Decarbonization." *Nature Energy* 2: 918–922. Available at https://doi.org/10.1038/s41560-017-0025-8.

Murray, B., and P. Maniloff. 2015. "Why Have Greenhouse Emissions in RGGI States Declined? An Econometric Attribution to Economic, Energy Market, and Policy Factors." *Energy Economics* 51(C): 581–589. Available at https://doi.org/10.1016/j.eneco.2015.07.013.

Olson, M. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, Cambridge.

Pahle, M., D. Burtraw, C. Flachsland, N. Kelsey, E. Biber, J. Meckling, O. Edenhofer, and J. Zysman. 2017. "What Stands in the Way Becomes the Way: Sequencing in Climate Policy to Ratchet Up Stringency Over Time," RFF Report.

Perino, G. 2018. "New EU ETS Phase 4 Rules Temporarily Puncture Waterbed." *Nature Climate Change* 8(4): 262–264.

Portney, P.R. 1990. "Economics and the Clean Air Act." *Journal of Economic Perspectives* 4(4): 173–181.

Roberts, M., and M. Spence. 1976. "Effluent Charges and Licenses Under Uncertainty." *Journal of Public Economics* 5(3–4): 193–208.

US EPA. 2002. "An Evaluation of the South Coast Air Quality Management District's Regional Clean Air Incentives Market—Lessons in Environmental Markets and Innovation." U.S. Environmental Protection Agency Region 9: Air Program, San Francisco, CA (November).

Van den Bergh, K., Delarue, E., and D'haeseleer, W. 2013. "Impact of Renewables Deployment on the CO_2 Price and the CO_2 Emissions in the European Electricity Sector." *Energy Policy* 63(C): 1021–1031.