The Welfare Implications of Carbon Price Certainty

Joseph Aldy and Sarah Armitage*

Draft: December 14, 2019

1 Introduction

Policy uncertainty has a material impact on firm behavior. Numerous papers have established this fact, beginning with Arrow and Fisher 1974 and Bernanke 1983 and more recently in the empirical work of Rodrik 1991; Hassett and Metcalf 1999; and Baker, Bloom, and Davis 2016. Much of this literature has focused on the impacts of uncertainty over future policy design, such as how the central bank will set future monetary policy or what future tax rates will be. A special case of this policy uncertainty occurs when the policy framework is fully resolved, so firms know what the future policy will be — for example, the regulator has decided to price some technological externality, settled on a price or quantity instrument, and determined its stringency. However, the characteristics of certain policies may create further uncertainties in the market that are never resolved, as these uncertainties are inherent to the policy instrument itself. We examine the impact of this residual policy uncertainty on the relative efficiencies of price- versus quantity-based policies for correcting technological externalities, such as greenhouse gas emissions.

*Acknowledgements: Jerry Green, Robert Pindyck, Edward Glaeser, Samuel Hanson, Jim Stock, Joshua Pfeffer for helpful discussions. Also participants at AERE Summer Conference, Northeast Environmental Economics Workshop, and Harvard Kennedy School Energy Policy seminar. Kenneth Norris for research assistance. Alex Barron for providing data.

Much of the literature comparing cap-and-trade programs and emission taxes has focused on the regulator's information deficit. Firms typically know their marginal abatement costs with greater precision than the regulator, and firms may not have the incentive to reveal their true marginal abatement costs to the regulator (and to their competitors). However, under cap-and-trade programs, this information asymmetry affects not only the regulator, but also each firm with respect to other firms' marginal abatement costs. Consider the problem a firm faces in complying with an emission tax versus a cap-and-trade program. Under any policy instrument, firms must first resolve uncertainty surrounding policy design. Then, if the regulator chooses a tax, the firm learns the tax rate, identifies its abatement options, and complies with the policy by investing in abatement that equates marginal cost with the tax. If the regulator instead chooses a cap-and-trade program, the firm not only must identify its abatement options as in the tax case, but also must form expectations about the market-clearing price for emissions allowances which will then guide the firm's investment.

Our paper examines this residual uncertainty about allowance prices, inherent to cap-and-trade programs, which increases the risk that firms may not equate their marginal abatement costs. Firms may err in their allowance price forecasts and make investments that would appear to be optimal ex ante given their expectations, but, ex post, are recognized as having been too high or too low. Such forecast errors may reflect different expectations about (1) abatement technology costs; (2) economic output; (3) overlapping public policies that may restrict abatement decisions and influence the clearing price in allowance markets. Given the common role of banking in cap-and-trade programs, these expectations would need to reflect a dynamic assessment of these factors. This residual uncertainty coupled with irreversible investments can increase the welfare costs of choosing a quantity-based instrument, such as cap-and-trade, relative to a carbon tax. Given the potential scale of tradable allowance markets, such forecast errors under cap-and-trade could be economically significant.

We build a theoretical model that relates these forecast errors to standard models of cap-and-trade programs and examines how firms make dynamic abatement decisions in the

presence of price uncertainty. One reason why existing literature has not focused on firm-level uncertainty over market-clearing permit prices is that intertemporal arbitrage requires that equilibrium prices are expected to rise at the rate of interest, following a Hotelling rule, whenever banking (and borrowing) of permits is permitted. However, observed volatility in allowance prices suggests that prices are continually jumping from one Hotelling trajectory to another, as shocks to price *levels* are realized. Our model first shows that price trajectories still follow a Hotelling rule in expectation, but the overall price process is determined by volatility as well as this expected drift. This volatility creates potentially significant forecast errors, as the realized trajectory of prices does not follow a smooth Hotelling trajectory.

The fact that firms often make dynamic decisions to invest in long-lived abatement capital magnifies the impact of these forecast errors. By deriving the necessary optimality conditions for firm investment in abatement, we show that firms not only set their marginal abatement costs equal to the discounted expected stream of allowance prices, as in standard models of abatement as a variable input, but also seek to smooth their rate of investment over time. Given uncertainty over allowance prices, the value of smoothing also depends on the extent of price volatility. This finding introduces the possibility that firms might err not only in their expectations of the realized allowance price, but also in their beliefs about the overall price process, as the volatility parameter depends on many different random processes, only some of which are observable to individual firms.

To understand the potential magnitude of these forecast errors, we turn to simulations calibrated to the performance of actual and modeled climate policies. We first simulate stochastic allowance price trajectories, with drift and volatility parameters estimated from historical prices from the European Union's Emissions Trading System (EU ETS). We model the corresponding abatement investment response by firms using an abatement cost function estimated from models of U.S. carbon pricing policy. We then identify the smoothly increasing "Hotelling" price trajectories that would achieve the same overall emissions reductions as each of the stochastic price simulations, and compare the total resource cost required to

achieve a particular level of emissions reductions under the two pricing scenarios. We show that the "Hotelling" scenario serves as a lower bound on total resource costs for achieving a given level of emissions control, by eliminating forecast errors on the part of regulated firms. We also show that it is possible to reduce the magnitude of these forecast errors through price collars.

To examine the welfare impacts of these forecast errors, we develop a modified version of Weitzman's canonical prices versus quantities comparison. In our version, quantity orders are not imposed directly on individual firms, as in Weitzman 1974 and much of the subsequent literature, but are instead transmitted to firms through some market-clearing price, as in modern allowance and credit trading markets. We show that firm-level forecast errors push the regulator to favor price-based instruments over quantity-based instruments, with the relative benefit of price instruments increasing in the variance of the forecast error term. Likewise, we also discuss how evaluations of price ceilings and floors in cap-and-trade programs, as hybrid policies between purely price-based and purely quantity-based instruments, should take into account the effectiveness of price collars at reducing the costs of forecast errors.

Section 2 presents empirical facts about real-world allowance and credit trading programs that together suggest that this residual policy uncertainty is likely to be economically meaningful. Section 3 develops our theoretical model of firm abatement decisions given investment in long-lived abatement capital and uncertainty over future allowance prices. Section 4 calibrates a simulation model to illustrate the potential magnitude of price forecast errors in these markets, using parameters calibrated to historical and modeled climate policies. Section 5 discusses how we can interpret these forecast errors in a Weitzman-style prices versus quantities welfare comparison. Finally, Section 6 concludes and offers directions for future research.

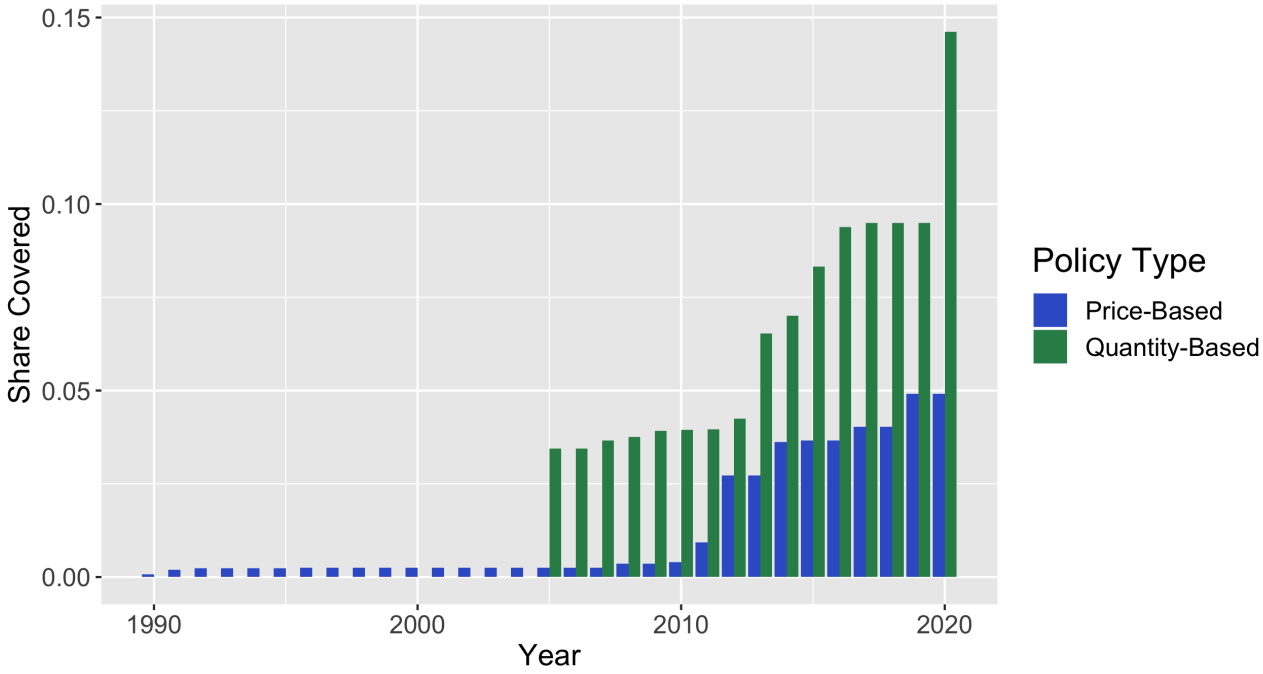


Figure 1: Global Carbon Emissions Covered by Policy (Ramstein et al. 2018).

2 Background

Three empirical observations motivate this paper. First, policymakers at the supranational, national, and sub-national levels are moving forward with carbon pricing policies. The World Bank recently estimated that nearly 20 percent of the world's carbon dioxide emissions are covered (or will soon be covered) by some carbon pricing policy (see Figure 1). While numerous hybrid policies exist, the majority of existing or planned policies are emissions trading programs rooted in quantity targets, rather than tax policies. Given the economic scale of the task of decarbonizing the global energy system in line with the goals established in the 2015 Paris Agreement, these policies will need to drive trillions of dollars of investment. Thus, poorly designed policy strategies may risk imposing significant costs.

Beyond climate policy, regulators have also exhibited strong revealed preference for using allowance or credit trading programs to correct other environmental externalities. In the United States, for example, virtually all energy produced or consumed is subject to some sort allowance or credit trading program. These programs range from state-level trading

programs for renewable energy credits as part of Renewable Portfolio Standards, to the nationwide sulfur and benzene credit trading program for gasoline, to the Acid Rain Program and the NO_x Budget Trading Program for controlling SO_2 and NO_x emissions, respectively. Understanding how well these programs have functioned is therefore critical for controlling environmental externalities as efficiently as possible.

The second key empirical observation relates to the nature of uncertainties under emission taxes versus allowance trading programs. Under any policy for controlling pollution, firms must first resolve uncertainty surrounding policy design. Will the regulator choose a cap-and-trade program, a tax, some variant that draws from each, or another policy approach? Then, after these questions about policy design are resolved, firms may continue to face considerable uncertainty as they make their compliance decisions, depending on the characteristics of the policy instrument selected. If the regulator chooses a tax, then the firm learns the tax rate, identifies its abatement options, and then complies with the policy by investing in abatement that equates marginal cost with the tax. If the regulator chooses a cap-and-trade program, then the firm still has to identify its abatement options as it did in the tax case. However, under cap-and-trade, a regulated firm also faces an uncertain market-clearing price for allowances. The firm therefore has to form expectations about all other firms' abatement functions, since these will translate into the supply and demand for emissions in the secondary tradable allowance market and thus the allowance price with which the firm would use as a guide for its investment.

This residual uncertainty inherent to cap-and-trade increases the risk that firms may not equate their marginal abatement costs. Firms may err in their allowance price forecasts and make investments that would appear to be optimal ex ante given their expectations, but, ex post, are recognized as having been too high or too low. Such forecast errors may reflect different expectations about (1) abatement technology costs; (2) economic output; (3) overlapping public policies that may restrict abatement decisions and influence the clearing price in allowance markets. Given the common role of banking in cap-and-trade programs,

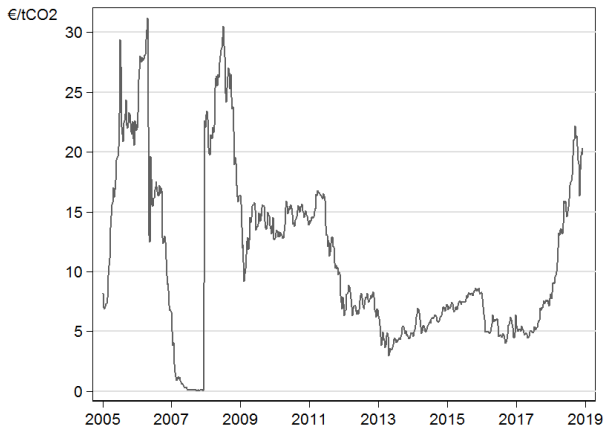
these expectations would need to reflect a dynamic assessment of these factors. When firms cannot easily reverse their abatement investments to correct their forecast errors, this residual uncertainty can increase the welfare costs of choosing a quantity-based instrument, such as cap-and-trade, relative to a carbon tax.

The third and final empirical observation relates to characteristics of real-world cap-and-trade programs which suggest that the magnitude of firm-level forecast errors may be large. For one, many of the most cost-effective abatement options in these markets are long-lived capital investments that are (to first order) irreversible. For example, under the Acid Rain Program, common abatement decisions were installing flue-gas desulfurization systems ("scrubbers") in power plants or retrofitting plants to accept low-sulfur coal, both of which were capital-intensive, largely permanent decisions (Cicala 2015). Likewise, in the NO_x Budget Trading Program, one of the most effective options for reducing regulated emissions was installing Selective Catalytic Reduction (SCR), which had the technical capacity to reduce emissions by up to 90% at the high capital cost of \$70-80/kW. In carbon trading markets, cost-effective abatement options might include investing in large-scale solar or wind generation, installing carbon capture and storage at power generation or industrial facilities, or retrofitting the built environment. Under other policies with credit trading programs, such as U.S. fuel efficiency standards for cars and trucks, firms might introduce new product models or redesign existing ones. In each of these cases, firms would need to make decisions about their abatement investments before information about allowance prices is fully realized. The temporal lag between investment decisions and market-clearing creates the possibility for significant forecast errors and undermines the standard assumption in literature on externality trading markets, that firms perfectly equate their marginal abatement costs with equilibrium allowance or credit prices.

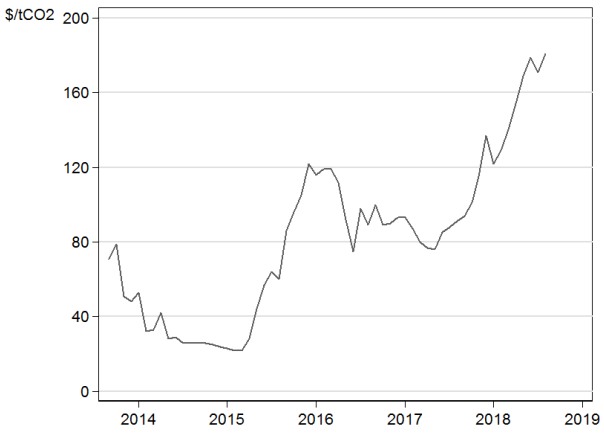
Furthermore, historical prices in allowance and credit trading markets have been extremely volatile, which also suggests that uncertainty about future prices may be substantial. Figure 2 illustrates the extent of historical price volatility in four different allowance or

credit trading programs. Under the EU ETS, market-clearing prices for carbon have ranged from over €30 per ton of CO_2 to under €5 per ton in the span of five years. Likewise, prices for SO_2 allowances under the U.S. Acid Rain Program have ranged from $1600 per ton to $100 per ton in a five-year period, and prices for NO_x allowances have ranged from $4,500 per ton to $800 per ton under the NO_x Budget Trading Program. Note that the observed volatility in these markets is not simply a function of the volatility of the underlying energy commodities, whether oil, natural gas, or coal. Indeed, the allowance price volatility observed under the EU ETS and the U.S. Acid Rain Program exceeds the volatility of oil or natural gas futures prices over comparable periods (see Figure 4). In all four markets, the volatility of allowance or credit prices exceeds the volatility of S&P 500 index prices over a comparable period.

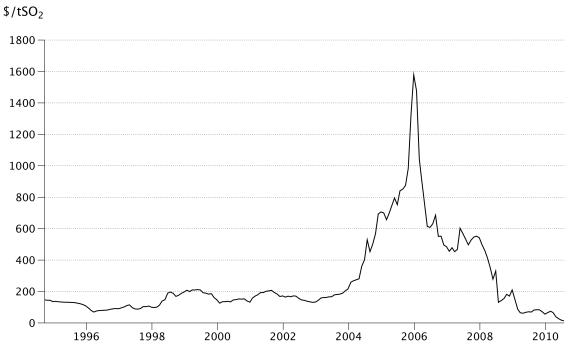
While it is possible that financial instruments would help regulated firms to mitigate uncertainty associated with volatile allowance and credit prices, evidence on hedging decisions more generally suggests that firms are likely to hedge incompletely, if at all. In studying hedging of input fuel prices by U.S. airlines, Rampini, Sufi, and Viswanathan 2014 find that the airlines in their sample hedge only 20% of expected next-year jet fuel expenses – despite the fact that financial instruments are widely available in this market and jet fuel represents a substantial and highly volatile operating expense for these firms. The authors attribute this imperfect hedging partly to firm financial constraints, which would certainly be relevant in our setting. Moreover, firms may also be unable to hedge their full exposure to uncertain allowance prices since the total quantity of allowances demanded depends on both the uncertain future price and potential additional uncertainties around future abatement cost (Gopinath). Finally, since allowance and credit trading markets are created virtually overnight through regulation, there is considerable uncertainty associated with the start-up of these markets which may reduce the availability of financial instruments in their early phases. To illustrate the potential magnitude of this start-up uncertainty, consider that proposals for U.S. economy-wide carbon cap-and-trade programs under the Kyoto Protocol



(a) Emissions Trading System CO_2 Allowance Price (EU)



(b) Low Carbon Fuel Standard Allowance Price (California)



(c) SO_2 Allowance Price (U.S.)

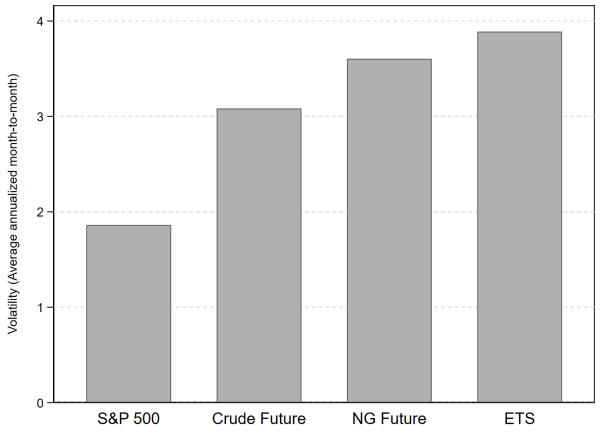


(d) NO_x Allowance Price (U.S.)

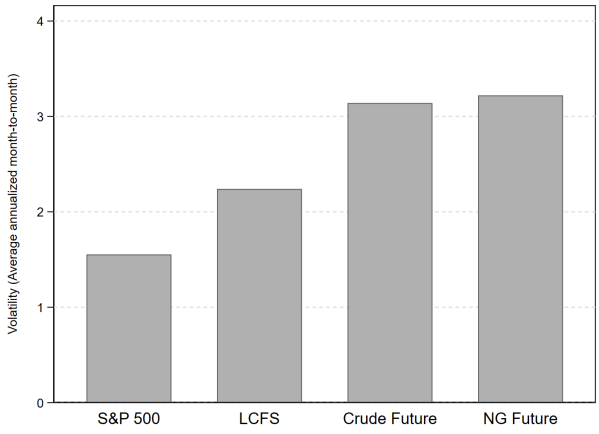
Figure 2: Historical Prices in Allowance and Credit Trading Markets

and the Waxman-Markey Bill would have created multi-billion dollar markets in their first year of operation, without any historical data on market performance to guide the supply of financial instruments. Even regulators designing these programs predicted a wide range of initial market values through their modeling scenarios, ranging from $177 to $683 billion under the Kyoto Protocol and from $54 to $254 billion under the Waxman-Markey Bill (in 2018 $). .

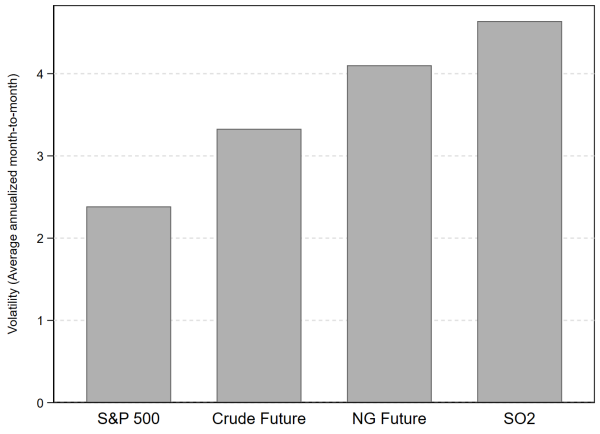
These empirical observations about the functioning of real-world allowance and credit trading programs suggest that the residual uncertainty inherent to these programs could have a significant impact on the overall cost of pollution reduction. Given the potential



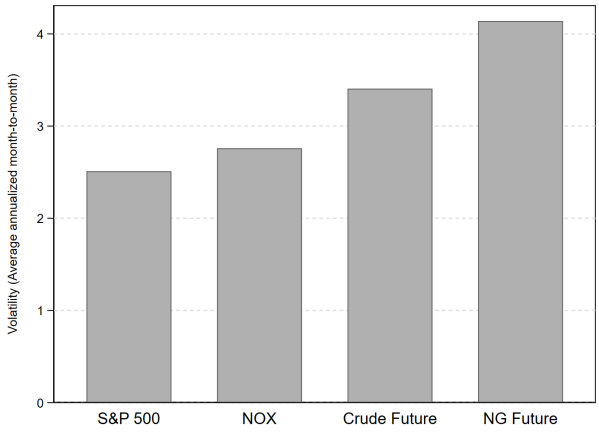
(a) Emissions Trading System CO_2 Allowance Price Volatility (EU)



(b) Low Carbon Fuel Standard Allowance Price Volatility (California)



(c) SO_2 Allowance Price Volatility (U.S.)



(d) NO_x Allowance Price Volatility (U.S.)

Figure 3: Average Annualized Month-to-Month Price Volatility in Allowance and Credit Trading Markets

scale of tradable allowance markets, as policymakers move forward with climate policies and continue to turn to these policy instruments for mitigating other types of environmental externalities, such forecast errors under cap-and-trade could be economically meaningful.

3 Existing Literature

Empirical literature on the functioning of real-world cap-and-trade markets has identified inefficiencies that are not easily explained by existing models of allowance trading. Carlson

et al. 2000 find that more than half of the units operating in Phase I of the SO_2 cap-and-trade program failed to minimize costs during at least part of the study period. Likewise, both Fowlie 2010 and Cicala 2015 show that deregulated firms may have underinvested in capital-intensive compliance strategies for SO_2 and NO_x cap-and-trade programs, combined with evidence of overinvestment by regulated firms. Overlapping policies may result in further deviations from cost minimization by increasing uncertainty in cap-and-trade markets. Frey 2013 finds, for example, that overlapping state regulations limited the cost-effectiveness of Phase I of the Acid Rain Program, by inducing higher-cost units to invest in scrubbers. Under an allowance trading program, the impact of such state policies would not only affect the units directly subject to state regulation, but also all other units in the allowance trading market. Uncertainty about the impact of overlapping policies translates into uncertainty about the net quantity limit in the allowance trading program, which then translates into further uncertainty over the market-clearing price that should guide compliance investments. In contrast, under a price-based program, the presence of overlapping policies does not affect the investment decisions of other firms in the market.

Our paper seeks to develop existing theories of cap-and-trade markets to account for these observations. Existing studies of allowance trading markets have generally adopted one of two assumptions which are incompatible with the actual functioning of these programs. On the one hand, papers in the spirit of Weitzman's canonical work on the relative advantage of price versus quantity advantage instruments have assumed that quantity orders are imposed directly on individual firms (Hoel and Karp 2002; Pizer 2002; Newell and Pizer 2003). These papers generally focus on uncertainties on the part of the regulator rather than on the part of regulated firms; one exception is Yohe 1978, which considers the possibility that firms may not know their own abatement cost functions with certainty, so quantity orders imposed directly on firms may not be achieved exactly. By contrast, in real-world allowance and credit trading markets, aggregate quantity orders are transmitted to firms through some (uncertain) market clearing price. One consequence of the complex price formation process is

that all types of shocks affect the equilibrium price, so regulated firms are not only impacted by direct shocks to their own abatement cost functions.

On the other hand, another body of existing literature that has explicitly modeled the market-clearing price associated with a given quantity limit has instead assumed that prices follow a Hotelling trajectory given banking and borrowing (Rubin 1996; Cronshaw and Kruse 1996; Newell, Pizer, and Zhang 2005). In these models, equilibrium allowance prices are pinned down by an inter-temporal arbitrage condition which requires prices to rise smoothly at the rate of interest. While Rubin achieves this Hotelling trajectory by assuming perfect certainty in the allowance trading market, Newell, Pizer, and Zhang examine how to maintain this smooth trajectory when shocks to abatement cost functions are realized each period. They show that banking and borrowing combined with a commitment on the part of the regulator to fix the last period expected price at a certain level will achieve this price trajectory. However, they do not explore how firm behavior would differ if allowance prices were allowed to remain stochastic, which is the focus of this paper. Real-world allowance prices evidently do not follow this smooth trajectory, perhaps suggesting that allowance prices are continually jumping from one Hotelling path to another as shocks are realized in the market, which we explore further below.

We therefore develop a model of cap-and-trade markets that integrates the two key empirical facts that market-clearing allowance prices are volatile and uncertain and that many abatement options are long-lived capital investments. Our approach combines three modeling strategies: Rubin's work on allowance trading markets with certainty and variable abatement decisions; Anderson, Kellogg, and Salant 2018 on dynamic investment decisions in natural resource extraction; and Pindyck 1980 on natural resource extraction under uncertainty. Our model recognizes that quantity orders are transmitted to individual firms through an uncertain market-clearing price; that firms may take into account price uncertainty in making abatement decisions; and that abatement decisions are linked over time.

The two papers that relate most closely to this paper in directly accounting for firm

investment decisions given allowance price uncertainty are Chao and Wilson 1993 and Zhao 2003. Chao and Wilson estimate the option value component of the SO_2 allowance price in the Acid Rain Program; this option value term reflects the fact that purchasing allowances may provide regulated firms with greater flexibility given uncertainty about future market conditions, as compared to investing in long-lived scrubbers. Then, Zhao 2003 considers how firm investment depends on uncertainty over the abatement cost function, where firms may invest in long-lived abatement capital or "technology" that reduces the cost of achieving a given level of abatement; in his model, actual compliance decisions are modeled as variable inputs. In comparing the impact of abatement cost uncertainty on both tradeable allowance programs and emission taxes, Zhao finds that tradeable allowance programs can be beneficial for responding to these uncertainties, as prices adjust to leave the overall quantity of variable abatement unaffected.

Our paper serves as a useful complement to this earlier work. Zhao's model considers how both firm-level and market-wide shocks affect firm investment in abatement technology but assumes that only market-wide shocks directly affect equilibrium prices or firms' marginal abatement costs. In contrast, our model allows price volatility to stem from any of the myriad random shocks affecting allowance trading markets, only some of which would affect firm abatement costs directly. Zhao's model also assumes a rational expectations equilibrium, whereas our paper explicitly considers the welfare consequences of forecast errors as part of the trade-off between price- and quantity-based policy instruments. Finally, unlike Zhao's model, our approach incorporates banking and borrowing of allowances, thereby allowing us to recover a version of the Hotelling rule for the expected growth in allowance prices that embeds existing models of allowance trading as a special case.

4 Theoretical Model of Abatement Investment

To build intuition for the impact of allowance price uncertainty, we develop a theoretical model of dynamic abatement investment decisions. This model focuses on the firm's cost-minimizing compliance strategy. We abstract from changes to the firm's production levels, though Montgomery 1972 shows how the cost of achieving a certain level of emissions may include both foregone profit due to deviations from unconstrained optimal production levels and the direct costs of installing abatement equipment. Our model shows that firms making abatement decisions will take into account estimated price volatility as part of their optimizing behavior. While the expected price trajectory is pinned down by a Hotelling rule, as in previous papers, the price volatility depends on a wide variety of random processes that affect this market. This volatility causes prices to jump from one Hotelling trajectory to another, thereby causing firms to err in their beliefs about the realized allowance price. Firms may also not know all of the random processes that determine price volatility, causing them to err further in their forecasts of the overall price process.

Our model integrates modeling strategies from three related papers. First, we adopt the basic model of an allowance trading market, with banking and borrowing of allowances across periods, developed in Rubin 1996. As with much of the standard literature on multi-period allowance trading markets, Rubin's model assumes no accumulation of abatement from one period to the next. We therefore build on Rubin's work by modeling long-lived investment as the key margin of firm decision-making around compliance. Our approach is similar in the spirit to Anderson, Kellogg, and Salant 2018, which derives a modified version of the Hotelling rule for natural resource extraction where the key margin of decision-making is investment in well drilling rather than production from already-drilled wells. Lastly, we also recognize that firms face a stochastic price process for allowances. We allow firms to take into account price volatility in their optimization decisions, following methods used in Pindyck 1980 for natural resource extraction under uncertainty.

4.1 Model Set-Up

Assume that a representative firm chooses some optimal rate of abatement investment $A(t)$ and allowance purchases $Y(t)$ at each instant, where positive $Y(t)$ corresponds to net allowances purchases and negative $Y(t)$ corresponds to net allowances sales. The investment cost function is given by $\psi(\cdot)$ and the instantaneous price of allowances is $P(t)$. The firm discounts future costs and benefits at some exogenous discount rate r. The firm accumulates investment in abatement over time, and the total stock of abatement capital is given by $K(t)$; this stock depreciates at rate δ. The firm is also able to bank unused allowances over time, where $B(t)$ gives the total size of the firm's bank at time t. The firm's baseline emissions rate is $\bar{E}$, of which $\bar{G}$ is grandfathered into the program by the regulator. Therefore, at each instant, the firm's compliance requirement is $\bar{E} - \bar{G} - K(t)$, which may be positive or negative and which may result in the firm buying or selling allowances or adding or removing allowances from its bank. Finally, following Anderson, Kellogg, and Salant 2018, we define $R(t)$ as the remaining abatement opportunities available to the firm, perhaps due to technological limitations on the firm's ability to reduce its emissions rate beyond a particular amount.

Therefore, at the start of the allowance trading market, the firm's problem is given by:

$$\max_{A,Y} \mathrm{E}_0[\int_0^\infty e^{-rt}\{-\psi(A(t)) - P(t)Y(t)\}dt] \tag{1}$$

subject to:

$$\dot{K} = A(t) - \delta K(t) \tag{2}$$

$$A(t) \geq 0,\ K_0 \text{ given} \tag{3}$$

$$\dot{B} = K(t) + Y(t) + \bar{G} - \bar{E} \tag{4}$$

$$B(t) \geq 0,\ B_0 = 0 \tag{5}$$

$$\dot{R} = -A(t) + \delta K(t) \tag{6}$$

$$R(t) = \bar{E} - K(t) \geq 0, \ R_0 \text{ given} \tag{7}$$

$$dP = \alpha P dt + \sigma P dz \text{ , where } dz \text{ is an increment of a Weiner process} \tag{8}$$

Let us discuss each of these constraints in turn. Equation 2 describes the law of motion for abatement capital; at each instant, investment at rate $A(t)$ adds to the stock, while depreciation at rate $\delta K(t)$ depletes the stock. Equation 3 sets the initial stock of abatement capital and constrains the rate of investment to be weakly positive, meaning that firms cannot reverse investments once they have occurred, besides waiting for depreciation to run its course. Likewise, Equation 4 describes the law of motion for the allowance bank. Additional allowances are added to the bank at the rate of allowances purchased $Y(t)$, net of allowances used for compliance $\bar{E} - \bar{G} - K(t)$. Equation 5 sets the initial level of the bank and constrains the bank to be weakly positive (a no-borrowing constraint). This no-borrowing assumption is easily relaxed, as we discuss below.

Equation 6 gives the law of motion of the stock of abatement opportunities, where current investment results in the loss of future opportunities, but depreciation of installed abatement capital creates additional opportunities. Here we define the stock of abatement opportunities as total baseline emissions less installed abatement capital, but we could assume technological limitations on the firm's ability to abate some of its baseline emissions without altering any of the major results below. Finally, we assume that prices follow a Geometric Brownian Motion process, with drift parameter α and volatility parameter σ, following Equation 8. We discuss in detail below the implications of assuming an exogenous price process in deriving the firm's optimal behavior.

To solve for the firm's necessary optimality conditions, we follow Pindyck 1980 and Karp and Traeger 2013 and define the flow profit function Π and the optimal value function J:

$$\Pi(t) = e^{-rt}\{-\psi(A(t)) - P(t)Y(t)\}$$

$$J(t,P,K,B,R) = \max_{A,Y} \mathrm{E}_t[\int_t^{\infty} e^{-r\tau}\{-\psi(A(\tau)) - P(\tau)Y(\tau)\}d\tau]$$

We then write the firm's problem as a dynamic programming equation in continuous time and expand using Ito's Lemma:[1]

$$0 = \max_{A,Y} \Pi(t) + \frac{1}{dt}\mathrm{E}_t[d(J)]$$

$$\begin{aligned} 0 = \max_{A,Y} \Pi(t) + J_t + J_B(K + Y + \bar{G} - \bar{E}) + J_K(A - \delta K) + J_R(-A + \delta K) \\ + J_P\alpha P + \frac{1}{2}J_{PP}\sigma^2 P^2 \end{aligned} \tag{9}$$

For simplicity, we assume the firm chooses an interior solution for both $A(t)^*$ and $Y(t)^*$, and the no-borrowing constraint is not immediately binding.[2] In the case of $Y(t)^*$, Rubin 1996 notes that if we wish to consider arbitrary price paths in this market, we would need to apply a technical condition to ensure that firms do not seek to buy or sell an unbounded number of allowances, since the firm's maximization problem is linear in the number of allowances purchased or sold. By applying the first-order condition directly in this derivation, we are assuming that the process of price formation is such that no firm finds it optimal to buy or sell an unbounded number of allowances in equilibrium. In the case of $A(t)^*$, our optimality condition dictates firm investment conditional on some interior solution $A(t)^* > 0$; it is possible that the first-order condition does not hold with equality for any positive value of $A(t)^*$ in which case the firm will be bound by the irreversibility constraint $A(t) \geq 0$ and set $A(t)^* = 0$.

With these assumptions in place, taking the first-order conditions with respect to the choice variables $A(t)$ and $Y(t)$ yields:

$$\frac{\partial \Pi}{\partial A(t)} + J_K - J_R = 0$$

[1] As in Karp and Traeger 2013, we are able to eliminate the expectation operator here since we are taking the expectation conditional on information at time t, so all functions are measurable at time t.

[2] For additional discussion of cases where the no-borrowing constraint binds, see Rubin 1996.

$$\frac{\partial \Pi}{\partial Y(t)} + J_B = 0$$

which can be rewritten as:

$$\psi'(A(t))e^{-rt} = J_K - J_R \tag{10}$$

$$P(t)e^{-rt} = J_B \tag{11}$$

4.2 Optimal Behavior in the Allowance Market

First, we focus on the optimality conditions governing firm behavior in the allowance market. We evaluate Equation 9 at the optimal values of $Y(t)^*$ and $A(t)^*$ and differentiate with respect to B:

$$\begin{aligned} 0 =& \frac{\partial \Pi}{\partial B} + J_{tB} + J_{BB}(K(t) + Y(t) + \bar{G} - \bar{E}) + J_{KB}(A(t) - \delta K(t)) + J_{RB}(-A(t) + \delta K(t)) \\ &+ J_{PB}\alpha P(t) + \frac{1}{2} J_{PPB}\sigma^2 P(t)^2 \end{aligned}$$

Applying Ito's Lemma allows us to rewrite this expression as:

$$0 = \frac{\partial \Pi}{\partial B} + \frac{1}{dt}\mathrm{E}_t[d(J_B)]$$

Since $\frac{\partial \Pi}{\partial B} = 0$, this expression simplifies to:

$$0 = \frac{1}{dt}\mathrm{E}_t[d(J_B)]$$

Substituting in the firm's first-order condition for $Y(t)$ (Equation 11) and expanding the differential operator gives:

$$0 = \frac{1}{dt}\mathrm{E}_t[d(P(t)e^{-rt})] = -rP(t)e^{-rt} + e^{-rt}\frac{1}{dt}\mathrm{E}_t[d(P(t))]$$

Rearranging terms gives:

$$rP(t) = \frac{1}{dt}\mathrm{E}_t[dP(t)] \tag{12}$$

This result recovers the standard Hotelling rule for the expected growth in equilibrium allowance prices – namely, that prices are expected to increase over time at the discount rate. It is useful to compare this result to the analogous finding in Rubin 1996, which assumes both perfect certainty about the future allowance market and no accumulation of abatement investment. When the no-borrowing constraint does not bind in his model, Rubin recovers the following rule for price dynamics:

$$rP(t) = \dot{P(t)} \tag{13}$$

In our model, Rubin's result holds in expectation. This finding then allows us to endogenize the drift parameter in our Geometric Brownian Motion process for P, as the Hotelling rule pins down that drift parameter α is equal to discount rate r in equilibrium.

4.3 Optimal Abatement Investment

Next, we examine optimal firm investment in abatement. We follow the same general steps described above; we first differentiate Equation 9 with respect to $K(t)$ and then apply Ito's Lemma, substitute the first-order conditions for $A(t)$ and $Y(t)$ in the resulting expression, and expand the remaining differential operator. This procedure yields the following condition:

$$(\delta + r)\psi'(A^*) = P + \frac{1}{dt}\mathrm{E}_t[d(\psi'(A^*))] \tag{14}$$

Note that $A(t)^*$ is a function of the firm's state variables B, K, and R, as well as price P. Therefore, we can expand the differential operator on the right-hand side of Equation 14

using the following identities:

$$\begin{aligned}\frac{1}{dt}\mathrm{E}_t[dA^*] &= \frac{1}{dt}\mathrm{E}_t[A_K dK + A_B dB + A_R dR + A_P dP + \frac{1}{2}A_{PP}dP^2] \\ &= A_K\dot{K} + A_B\dot{B} + A_R\dot{R} + A_P\alpha P + \frac{1}{2}A_{PP}\sigma^2 P^2\end{aligned} \tag{15}$$

$$\begin{aligned}\frac{1}{dt}\mathrm{E}_t[(dA^*)^2] &= \frac{1}{dt}\mathrm{E}_t[(A_K dK + A_B dB + A_R dR + A_p dP + \frac{1}{2}A_{pp}dP^2)^2] \\ &= A_P\sigma^2 P^2\end{aligned} \tag{16}$$

where the second set of equalities follows from eliminating all higher-order expressions of dt that vanish as $dt \to 0$.

Our optimality condition then becomes:

$$\begin{aligned}\psi(A^*) &= \frac{1}{r+\delta}\{P + \psi''(A^*)\frac{1}{dt}\mathrm{E}_t[dA^*] + \frac{1}{2}\psi'''(A^*)\frac{1}{dt}\mathrm{E}_t[(dA^*)^2]\} \\ &= \frac{1}{r+\delta}\{P + \psi''(A^*)\frac{1}{dt}\mathrm{E}_t[dA^*] + \frac{1}{2}\psi'''(A^*)A_p^2\sigma^2 P^2\}\end{aligned} \tag{17}$$

where the expression for $\frac{1}{dt}\mathrm{E}_t[dA^*]$ is given above.

Given an interior solution for the optimal rate of abatement investment, we find that the firm not only sets its marginal abatement cost equal to the discounted stream of avoided allowance payments, but also takes into account additional terms that capture the value of smoothing its rate of investment over time. This value of smoothing depends, in turn, on the extent of price volatility in the market, captured by the parameter σ. The value of smoothing also depends on both the convexity of the abatement cost function and of the marginal abatement cost function. The latter relationship may be understood as analogous to "prudence" terms in lifetime consumption-savings models in macroeconomics. In these models, convexity of marginal utility of consumption induces individuals to engage in precautionary savings in the presence of uncertainty. In our context, we can understand increased abatement investment as analogous to precautionary savings in the face of uncertain allowance prices, which occurs whenever the marginal abatement cost function is concave (such that

$-\psi'(\cdot)$ is convex).[3]

We can lend further interpretation to this optimality condition by considering what happens when we set price volatility to zero or instead model abatement decisions as variable. First, if we set price volatility to zero, we effectively assume that allowance prices will follow the same Hotelling trajectory indefinitely. In that case, we recover the following variant of Equation 17:

$$\psi'(A^*) = \frac{1}{r+\delta}\{P + \psi''(A^*)\dot{A}\} \tag{18}$$

In this modified optimality condition, the firm still wishes to smooth abatement investment over time, but now the time derivative of A^* is known with certainty.[4] The "prudence" term, which depends linearly on price volatility, also vanishes.

Likewise, it is also instructive to compare our result in Equation 17 to a model where abatement is a variable decision, rather than a dynamic investment. We could recover this scenario as a special case of the model presented above by assuming full depreciation ($\delta = 1$). In this case, the firm's first-order conditions become:

$$\frac{\partial \Pi}{\partial Y} + J_B = 0$$

$$\frac{\partial \Pi}{\partial A} + J_B = 0$$

which can be rewritten as:

$$e^{-rt}\frac{\partial \psi}{\partial A} + J_B = 0$$

$$e^{-rt}P + J_B = 0$$

Rearranging terms and combining these expressions immediately yields:

$$\psi'(A(t)) = P \tag{19}$$

[3]Provided that optimal abatement is decreasing in allowance price, or $A_P \leq 0$.

[4]This result has parallels to the modified Hotelling rule for oil drilling in Anderson, Kellogg, and Salant 2018, which models firms making investment decisions as if future prices were certain.

Here the firm sets the rate of abatement equal to the instantaneous price of allowances. This result matches the abatement optimality condition in Rubin 1996.

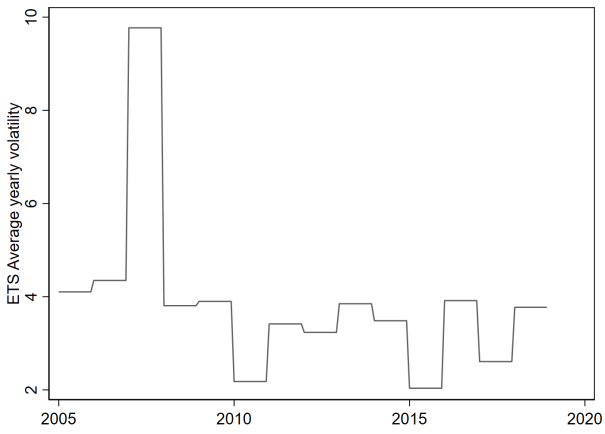
4.4 Discussion

Our optimality conditions show that allowance trading activity, combined with banking and borrowing, pins down expected price growth in equilibrium. We can then interpret the volatility parameter as causing the price process to jump continually from one expected Hotelling path to another. Where does this volatility come from? While much existing literature discusses the Hotelling rule and allowance price trajectories over time, considerably fewer papers have endogenized the overall price *level*. One exception is Stock, Irwin, and MCormack, Forthcoming, which determines that the volatility of RIN prices in the Renewable Fuel Standard credit trading market is largely determined by economic fundamentals such as the global prices of oil, corn, and soybeans – which are, at least to an extent, exogenous to decisions in the abatement market. From this perspective, assuming an exogenous volatility parameter may be a reasonable approximation. Additional sources of price volatility in these markets may also be, to first order, exogenous to the abatement behavior of market participants, including changes in economic output, changing expectations about future policy reforms, or unpredictable impacts of overlapping policies.

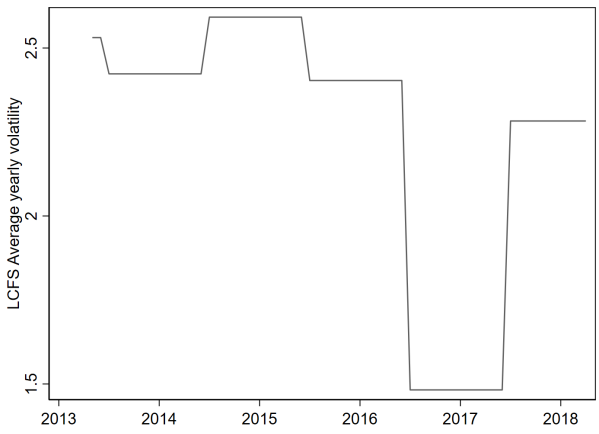
In a supplemental analysis (see Appendix), we also solved for the firm's necessary optimality conditions given an endogenous price process. In this extended version of the model, we model the market-clearing price as a flexible function of the overall emissions limit set by the regulator, the baseline level of emissions, total installed abatement capital, total accumulated allowance bank, and, critically, random shocks to both baseline emissions (e.g., due to changes in economic output or overlapping policies) and the abatement cost function (e.g., due to changes in the relative cost of high- versus low-emission fuels). This extended model therefore allows us to endogenize price volatility as a function of market fundamentals. Assuming a competitive market, in which firms do not internalize the impact of their deci-

sions on the market-clearing price or behave strategically, we find that the firm's optimality conditions bear close resemblance to those derived above. We again recover the Hotelling rule for the expected path of allowance prices, and the optimal level of abatement investment again comes from setting the marginal cost of abatement equal to the expected discounted stream of avoided allowgratance payments plus terms that capture the value of smoothing investment over time. However, now the smoothing terms depend on the volatility of the underlying random processes, instead of the volatility of the exogenous price process. Likewise, the "prudence" term is expanded to include cross-partial derivatives of the investment cost function between abatement level and the random shock to the cost function; this source of volatility – in contrast to shocks to baseline emissions – affects both the market-clearing price and, conditional on that price, the optimal level of abatement investment. However, our substantive conclusions remain unchanged after we endogenize price volatility in this manner, which lends additional support to our use of an exogenous price process in our baseline model.

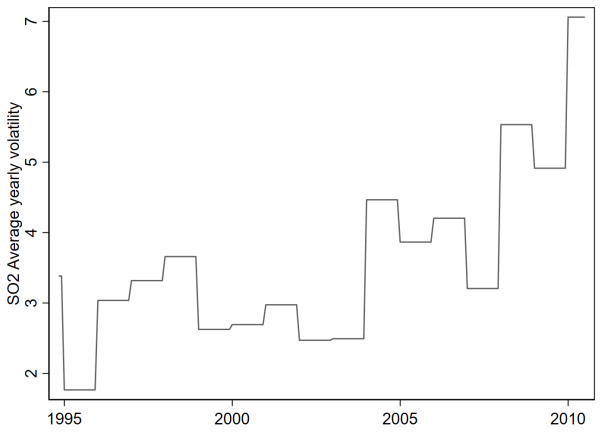
Given the significant price volatility observed in real-world allowance and credit trading markets, it seems likely that actual prices do indeed continually jump from one Hotelling path to another, as we have modeled here. As a consequence, firms are likely to experience significant forecast errors between the expected allowance price, against which firms make their abatement decisions, and the actual market-clearing price realized ex post. Moreover, firms may also error in their forecasts of the price *process*, in addition to specific realized prices. To the extent that the volatility parameter depends on many random variables, only some of which are known to an individual firm, each firm may not know the true value of σ when making its abatement decisions. Furthermore, firms may not be able to learn this parameter simply by observing the past history of prices, given that the extent of price volatility has varied significantly across different allowance and credit trading markets and over time (see Figure 3).



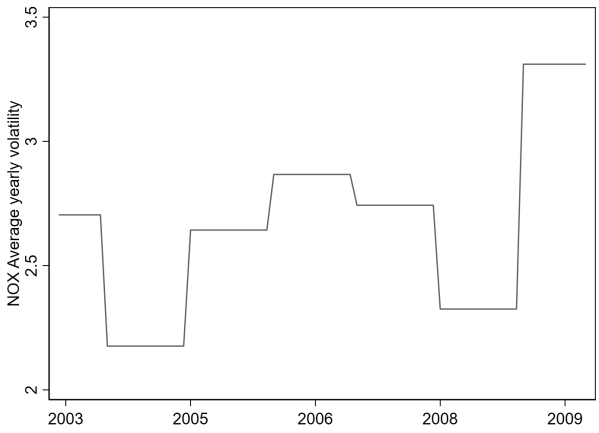
(a) Emissions Trading System CO_2 Allowance Price Volatility (EU)



(b) Low Carbon Fuel Standard Allowance Price Volatility (California)



(c) SO_2 Allowance Price Volatility (U.S.)



(d) NO_x Allowance Price Volatility (U.S.)

Figure 4: Yearly Price Volatility in Allowance and Credit Trading Markets

5 Simulations

To illustrate the potential magnitude of these forecast errors arising from price volatility in cap-and-trade markets, we have calibrated a simulation model based on the recent Stanford Energy Modeling Forum (EMF-32) study focused on U.S. carbon tax policies (Barron et al. 2018, McFarland et al. 2018). In developing this model, we have also incorporated the volatility of EU ETS prices as a real-world illustration of how uncertainties play out in tradable allowance markets.

We first simulate 10,000 different stochastic price trajectories using drift and volatility parameters estimated from the EU ETS. We then model the optimal abatement investment decisions of a representative firm faced with this stochastic price trajectory over ten periods, with prices leveling off indefinitely after the final period. We assume that the firm only knows the previous period's realized price when making its investment decisions, thereby allowing us to examine the impact of forecast errors around the realized price. We do not consider the impact of forecast errors related to the overall price process in these simulations.

Based on the total emissions reductions achieved from a given stochastic price trajectory, summed over all periods, we then calculate the "Hotelling" price trajectory that would achieve the same level of overall emissions reductions – that is, a single price that rises each period at the rate of interest, before leveling off indefinitely after the final period. This step allows us to compare the total resource costs required to achieve a given level of emissions reduction with and without firm-level uncertainty. We can also illustrate how price collars such as price floors and ceilings can reduce the cost inefficiencies resulting from this residual uncertainty.

5.1 Model Calibration

We first assume that allowance prices follow Geometric Brownian Motion and estimate the drift and volatility parameters associated with historical prices in the EU ETS. We focus on the periods since the end of the Phase I pilot program (2008-2018), as rules for intertemporal banking and borrowing changed between Phases I and II. We estimate these parameters using maximum likelihood; full details of the estimation procedure are provided in the appendix. We estimate an annual drift parameter of 0.065 and an annual volatility parameter of 0.3923; our drift parameter corresponds to 6.72% expected annual price growth, which is consistent with our earlier result that the drift parameter should be equivalent to the discount rate.

We then use results from the Stanford Energy Modeling Forum to calibrate the abatement

cost function. This dataset includes projected U.S. emissions reductions, relative to a baseline scenario, resulting from an emissions price set at either $25 or $50 in 2020 and increasing at an annual rate of either 1% or 5% to 2050. The dataset includes results from 10 models analyzing each of these four price scenarios.

We adopt the simplifying assumption that all abatement is long-lived and thus abatement in the current period persists into the next period, adjusting for depreciation; this assumption matches our theoretical modeling of abatement as durable capital stock. Likewise, we also assume that abatement investment in a given year becomes available for compliance in the following year. Based on these assumptions, we calculate the discounted value of the tax payment avoided through abatement investment in each period. By assuming that the investment cost function takes the form $\psi(A) = \phi A^2$, we use the firm's first-order conditions to set the marginal investment cost equal to the discounted stream of avoided tax payments. We calculate the depreciated sum of abatement and compare that to average emissions reductions observed in the modeling scenarios (relative to the baseline scenario). Setting these two values equal then allows use to estimate the abatement cost parameter ϕ.

We obtain an estimated parameter $\hat{\phi} = 6.97 \cdot 10^{-07}$ for abatement measured in metric tons of CO_2 reduced annually. Figure 5 shows the annual emissions reduction predicted by our calibrated model versus each of the ten modeling scenarios from the Stanford EMF-32 exercise.

5.2 Simulation Results

Given this calibration, Figure 6 illustrates the results of our simulations. The horizontal axis shows the different levels of total emissions reduction (in million metric tons of CO_2) resulting from different price simulations and ten years of firm investment in abatement. The vertical axis shows the total investment cost required to achieve those reductions. As is evident from the figures, the Hotelling price trajectories serve as a lower bound on the total reosurce costs for a wide range of abatement levels. The expected increase in costs

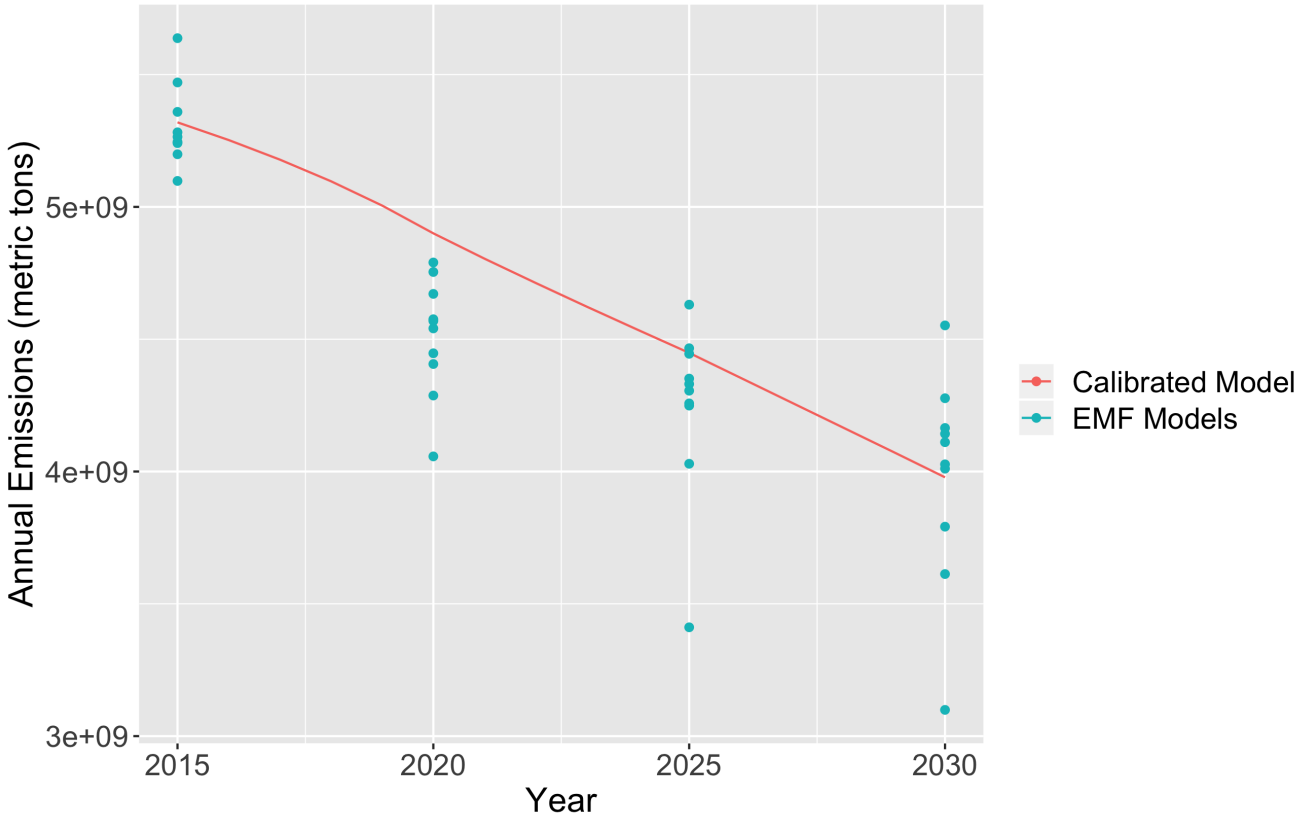


Figure 5: Annual Emissions from Calibrated Model versus EMF 32 Results

from stochastic prices increases in the magnitude of realized abatement.[5] These results illustrate why it is critical to account for the magnitude of forecast errors in considering the relative welfare gains from price or quantity instruments, as we discuss in more detail in the subsequent section. The presence of firm-level uncertainty over the market-clearing allowance price effectively shifts upwards the expected abatement cost function – that is, the expected total resource cost associated with achieving a given quantity of emissions reduction. Therefore, it is not sufficient to compare the abatement costs of a quantity order imposed directly on firms, on the one hand, with the abatement costs from setting marginal abatement cost equal to the regulated price, on the other hand (as in Pizer 2002, for example). Instead, the effective abatement cost function is itself a function of the type of policy instrument.

5.2.1 Price Floors and Price Ceilings

(MATERIAL ON ALLOWANCE PRICE CEILINGS AND FLOORS IS FORTHCOMING.)

6 Welfare

To understand the welfare consequences of firm uncertainty over market-clearing allowance prices, we turn to a modified version of the Weitzman 1974 derivation for the relative advantage of price-based instruments over quantity-based instruments. We first develop Weitzman's original derivation with multiple production units to account for the fact that aggregate quantity orders are transmitted through an equilibrium price. To be consistent with Weitzman's set-up, all randomness in the market-clearing price stems from stocks to each firm's marginal abatement costs; however, as we have discussed earlier, real-world price volatility stems from myriad other factors including changes in overall economic

[5]Note that we set an effective price ceiling at $1000 per ton when discretizing the state space to perform backwards induction; however, given our drift and volatility parameters and the number of periods considered, this upper bound affects less than 0.1% of simulated price trajectories.

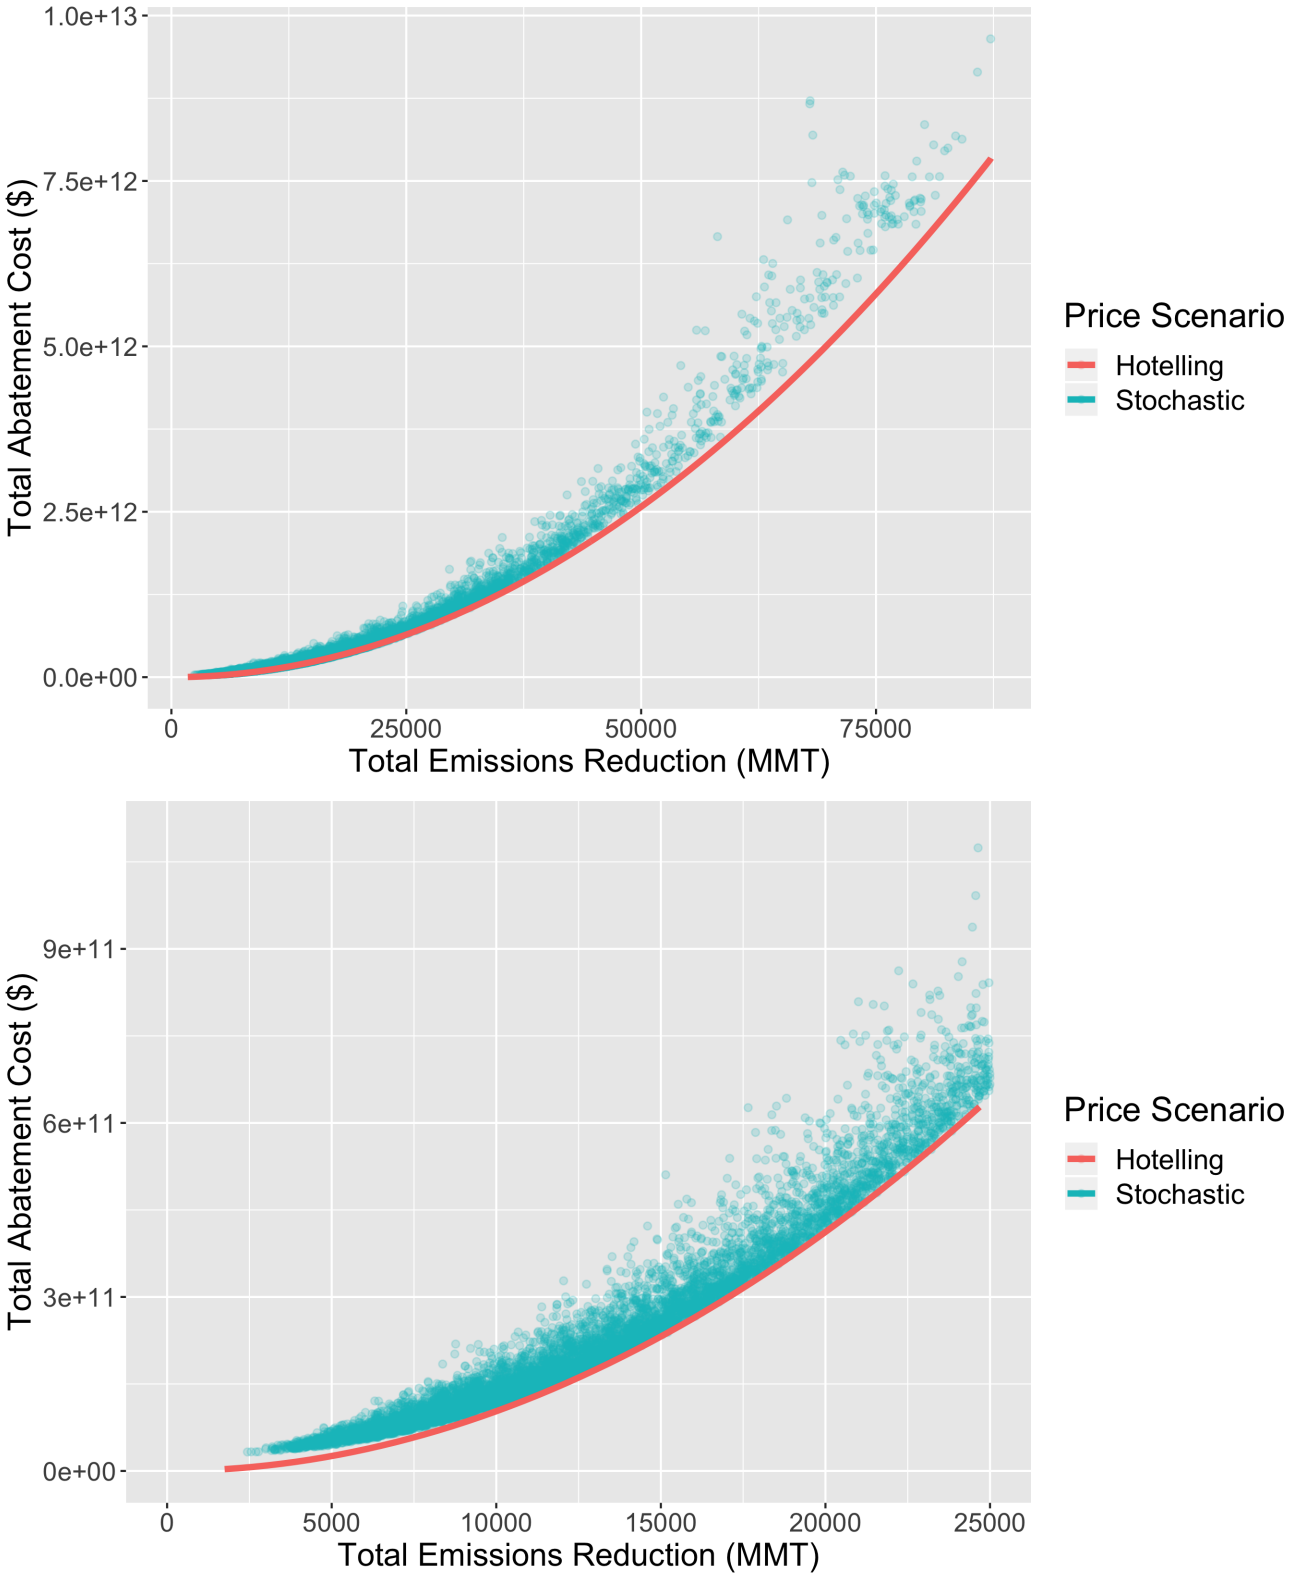


Figure 6: Total Emissions Reduction from 10 Years of Abatement Investment

output or the impact of overlapping policies.

We begin by defining the benefit function associated with reducing some pollutant and the cost function associated with abatement. Let $B(Q, \eta)$ be the benefits associated with producing quantity Q of the pollutant, and let η represent a random shock to the benefit function. Likewise, let $C_i(q_i, \theta_i)$ be the costs to firm i associated with producing quantity q_i of the pollutant, where θ_i represents a firm-specific random shock to the cost function. Therefore, the aggregate costs associated with pollution level Q are given by $\sum_i^N C_i(q_i, \theta_i)$, where $Q = \sum_i^N q_i$. Here we are assuming that the pollutant in question is uniformly mixed, such that only the total level of pollutant enters into the benefits function, not the identity of each polluting entity; this assumption reflects the characteristics of most greenhouse gas emissions but could be relaxed to model local pollutants. By contrast, the costs of abatement depend on the pollution level achieved by each individual firm.

Following Weitzman (1974), the regulator sets an optimal price order in the presence of uncertainty by solving the following maximization problem:

$$\max_p \mathrm{E}[B(\mathbf{h}(\mathbf{p}, \theta), \eta) - \sum_{i=1}^{N} C_i(h_i(p, \theta_i), \theta_i)] \tag{20}$$

Likewise, the regulator sets the optimal (aggregate) quantity order by solving the following maximization problem:

$$\max_Q \mathrm{E}[B(\mathbf{h}(p(Q, \theta), \theta), \eta) - \sum_{i=1}^{N} C_i(h_i(p(Q, \theta), \theta_i), \theta_i)] \tag{21}$$

Here $h_i(p, \theta_i)$ represents firm i's response function given price p and cost shock θ_i, which we describe further below.

Following Weitzman, we expand the cost functions by taking a second-order Taylor expansion about the quantity $\bar{q}_i$. We define each $\bar{q}_i$ as the level of the pollutant that sets expected benefits equal to expected costs for each individual firm. However, as we discuss further below, however, $\bar{q}_i$ does not necessarily represent the regulator's optimal quantity

target in this setting, since the regulator now chooses the aggregate quantity only.

$$\mathrm{E}[B_1^i(\bar{\mathbf{q}}, \eta)] = \mathrm{E}[C_1^i(\bar{q}_i, \theta_i)] \tag{22}$$

The Taylor expansion of each abatement cost function about $\bar{q}_i$ is then given by:

$$C_i(q_i, \theta_i) = a_i(\theta_i) + (C_i' + \alpha_i(\theta_i))(q_i - \bar{q}_i) + \frac{C_i''}{2}(q_i - \bar{q}_i)^2 \tag{23}$$

where C' represents the expected marginal abatement cost at $\bar{q}_i$ and C'' represents the slope of the marginal abatement cost function. As in Weitzman and much of the subsequent literature, we assume for tractability that the abatement cost function is quadratic or can be well approximated by a second-order Taylor expansion. $a_i(\theta_i)$ represents how the random cost shock θ_i affects the level of the cost function, and $\alpha_i(\theta_i)$ represents how the cost shock affects the slope.

For the benefits function, we take the Taylor expansion around $\bar{Q} = \sum_i^N \bar{q}_i$:

$$B(Q) = b(\eta) + (B' + \beta(\eta))(\sum_{i=1}^{N} q_i - \bar{q}_i) + \frac{B''}{2}(\sum_{i=1}^{N} q_i - \bar{q}_i)^2 \tag{24}$$

Here B' captures the expected marginal benefit at $\bar{Q}$, and B'' captures the slope of the marginal benefit function. $b(\eta)$ is the impact of the random shock η on the level of benefits, and $\beta(\eta)$ is the impact on the slope.

As in Weitzman, we assume without loss of generality that $\mathrm{E}[a_i(\theta_i)] = 0$, $\mathrm{E}[\alpha_i(\theta_i)] = 0$, $\mathrm{E}[b(\eta)] = 0$, and $\mathrm{E}[\beta(\eta)] = 0$. Note that this assumption does not preclude correlations across cost shocks to individual firms, so it is still possible that $\mathrm{E}[a_i(\theta_i)|a_j(\theta_j)] \neq 0$ or $\mathrm{E}[\alpha_i(\theta_i)|\alpha_j(\theta_j)] \neq 0$ for some i and j. We also assume that $\mathrm{E}[\beta(\eta) \cdot \alpha_i(\theta_i)] = 0$ for all i.

For the optimal price order, the derivation here closely follows Weitzman's derivation with multiple production units, except we constrain the regulated price to be the same across all units. We again find that the optimal price order $\tilde{p}$ is equal to $B' = C_i'$ for all i. (The full

derivation is provided in the appendix.) Assuming cost minimization, each firm will set its realized marginal cost function equal to this price, yielding the following firm-level response function:

$$h_i(\tilde{p}, \theta_i) = \tilde{q}_i = \bar{q}_i - \frac{\alpha_i(\theta_i)}{C_i''} \tag{25}$$

The aggregate quantity produced will then be:

$$\sum_{i=1}^{N} h_i(\tilde{p}, \theta_i) = \sum_{i=1}^{N} \bar{q}_i - \frac{\alpha_i(\theta_i)}{C_i''} = \bar{Q} - \sum_{i=1}^{N} \frac{\alpha_i(\theta_i)}{C_i''} \tag{26}$$

By contrast, for the optimal quantity order, we must solve for the market-clearing price such that the aggregate quantity order is achieved after the realization of all shocks in the market. Based on the set-up above and again assuming cost minimization, the equilibrium price associated with the optimal quantity order $\hat{Q}$ is given by:

$$\hat{p}(\hat{Q}, \theta) = \frac{\sum_{i=1}^{N} \frac{C_i' + \alpha_i(\theta_i)}{C_i''}}{\sum_{i=1}^{N} \frac{1}{C_i''}} \tag{27}$$

Consequently, firm i will produce $\hat{q}_i = \bar{q}_i + \frac{\sum_{j=1}^{N} \frac{\alpha_j(\theta_j)}{C_j''}}{C_i'' \sum_{j=1}^{N} \frac{1}{C_j''}} - \frac{\alpha_i(\theta_i)}{C_i''}$. The aggregate quantity produced will be $\hat{Q} = \bar{Q}$. (The full derivation is again given in the appendix.)

Finally, following Weitzman, we define the relative advantage of prices over quantities as the expected difference between net benefits from the optimal price order and net benefits from the optimal quantity order. That is:

$$\Delta = \mathrm{E}[B(\sum_{i=1}^{N} \tilde{q}_i, \eta) - \sum_{i=1}^{N} C_i(\tilde{q}_i, \theta_i)] - \mathrm{E}[B(\sum_{i=1}^{N} \hat{q}_i, \eta) - \sum_{i=1}^{N} C_i(\hat{q}_i, \theta_i)] \tag{28}$$

Substituting each firm's response to the price and quantity orders, respectively, we obtain

the following variant on Weitzman's original result with multiple production units:

$$\Delta = \mathrm{E}[\frac{B''}{2}(\sum_{i=1}^{N} \frac{\alpha_i(\theta_i)}{C_i''})^2 + \frac{1}{2\sum_{j=1}^{N} \frac{1}{C_j''}}(\sum_{i=1}^{N} \frac{\alpha_i(\theta_i)}{C_i''})^2] \tag{29}$$

Note that in this modified version of Weitzman's derivation, we cannot evaluate the costs associated with the quantity order for each firm separately, since the quantity produced by each firm depends on the shocks to all other firms' marginal abatement cost functions. Consequently, we must compare the slope of marginal benefits to an expression that combines the slopes of all marginal costs.

6.1 Expectation Errors

Because the market-clearing price associated with the regulator's quantity order depends on shocks to the marginal abatement cost functions of all firms, a given firm i may not know this price with certainty. As a result, we assume that firm i will set its marginal abatement cost function equal to the efficient price $\hat{p}$ associated with aggregate quantity order $\hat{Q}$, plus some expectation error ϵ_i. That is, $\mathrm{E}_i[\hat{p}(\hat{Q}, \theta)] = \hat{p}(\hat{Q}, \theta) + \epsilon_i$. The corresponding quantity response is given by:

$$h_i(\hat{p}, \theta_i) = \frac{\hat{p}(\hat{Q}, \theta) + \epsilon_i - C_i' - \alpha_i(\theta_i)}{C_i''} + \bar{q}_i \tag{30}$$

By contrast, under a price order, the price is set by regulation and known with certainty to all firms. Therefore, firms' response to a price order does not change from the version derived above.

Substituting the modified firm response to a quantity order into the expression for the relative advantage of prices over quantities, we obtain:

$$\Delta = \mathrm{E}[\frac{B''}{2}(\sum_{i=1}^{N} \frac{\alpha_i(\theta_i)}{C_i''})^2 + \frac{1}{2\sum_{j=1}^{N} \frac{1}{C_j''}}(\sum_{i=1} \frac{\alpha_i(\theta_i)}{C_i''})^2 + \sum_{i=1}^{N} \frac{\epsilon_i^2}{2C_i''}] \tag{31}$$

Here we place constraints on the expectation errors to ensure that the aggregate quantity

order is still met. We therefore require $\sum_{i=1}^{N} \frac{\epsilon_i}{C_i''}$, which ensures that firms' expectation errors will collectively cancel with each other in determining the overall quantity.

Equation 31 indicates that these firm-level forecast errors represent an advantage of price instruments relative to quantity instruments, with the relative advantage increasing in the variance of the error terms. One way to interpret this finding, in light of Weitzman's original result, stems from the fact that the regulator is no longer imposing quantity orders directly on individual firms. Instead, the regulator imposes an aggregate quantity order, which a market mechanism then translates into individual quantity orders through the market-clearing price and the magnitude of marginal cost shocks realized for each firm. Given firm-level expectation errors, the same relative advantages of price and quantity instruments still exist, but we must also consider the possibility that the aggregate quantity order is not distributed in the least cost manner across firms.

As Newell and Pizer 2003 notes, the original Weitzman derivation with multiple production units could also be understood as a representative firm subject to regulation over multiple periods. Likewise, we can also understand the welfare result derived above in a multi-period framework. In this case, the aggregate quantity order might represent some cap on a stock pollutant emitted over multiple periods, where regulated firms are free to bank and borrow allowances across periods. The price p now represents the discounted present value of future tax levels or allowance prices, and the forecast errors reflect incorrect estimates of that discounted sum because future shocks have not yet been realized.

7 Conclusion

This paper examines the impact of firm-level uncertainty over allowance prices in cap-and-trade markets – a form of residual uncertainty which is inherent to this type of policy instrument. Our theory model elucidates forecast errors that are not emphasized in the standard literature, both the difference between expected price and realized price and imperfect

information about the overall price process. Our simulations suggest that the magnitude of these forecast errors may be substantial in the context of climate policy, creating a wedge between the effective abatement cost function under price certainty versus in the presence of these forecast errors. Finally, we show that these forecast errors may be integrated into a standard prices versus quantities framework, by adding to the relative advantage of prices over quantity instruments.

In future research, we seek to decompose the extent of price volatility into that resulting from own abatement cost shocks versus that from shocks to the abatement costs of other firms, the impacts of overlapping policies, or changes to economic output. We also have ongoing work to incorporate an assessment of overlapping policies in our simulation model.

References

Stock, James H., Scott H. Irwin, and Kristen MCormack. Forthcoming. "The Price Of Biodiesel Rins and Economic Fundamentals". *American Journal of Agricultural Economics*.

Montgomery, W. David. 1972. "Markets in Licenses and Efficient Pollution Control Programs". *Journal of Economic Theory* 5:395–418.

Arrow, Kenneth J., and Anthony C. Fisher. 1974. "Environmental Preservation, Uncertainty, and Irreversibility". *The Quarterly Journal of Economics* 88 (2): 312–319.

Weitzman, Martin L. 1974. "Prices vs. Quantities". *The Review of Economic Studies* 41 (4): 477–491.

Yohe, Gary W. 1978. "Towards a General Comparison of Price Controls and Quantity Controls under Uncertainty". *The Review of Economic Studies* 45 (2): 229–238.

Pindyck, Robert S. 1980. "Uncertainty and Exhaustible Resource Markets". *Journal of Political Economy* 88 (6): 1203–1225.

Bernanke, Ben S. 1983. "Irreversibility, Uncertainty, and Cyclical Investment". *The Quarterly Journal of Economics* 98 (1): 85–106.

Rodrik, Dani. 1991. "Policy uncertainty and private investment in developing countries". *Journal of Development Economics* 36 (2): 229–242.

Chao, Hung-Po, and Robert Wilson. 1993. "Option Value of Emission Allowances". *Journal of Regulatory Economics* 5 (3): 233–249.

Cronshaw, Mark B., and Jamie Brown Kruse. 1996. "Regulated Firms in Pollution Permit Markets with Banking". *Journal of Regulatory Economics* 9:179–189.

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

Hassett, Kevin A., and Gilbert E. Metcalf. 1999. "Investment with Uncertain Tax Policy: Does Random Tax Policy Discourage Investment?" *The Economic Journal* 109:372–393.

Carlson, Curtis, et al. 2000. "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Resources for the Future Discussion Paper*, no. 98-44-REV.

Hoel, Michael, and Larry Karp. 2002. "Taxes versus quotas for a stock pollutant". *Resource and Energy Economics* 24:367–384.

Pizer, William A. 2002. "Combining price and quantity controls to mitigate global climate change". *Journal of Public Economics* 85 (3): 409–434.

Newell, Richard G., and William A. Pizer. 2003. "Regulating stock externalities under uncertainty". *Journal of Environmental Economics and Management* 45:416–432.

Zhao, Jinhua. 2003. "Irreversible abatement investment under cost uncertainties: tradable emission permits and emissions charges". *Journal of Public Economics* 87:2765–2789.

Newell, Richard G., William A. Pizer, and Jiangfeng Zhang. 2005. "Managing Permit Markets to Stabilize Prices". *Environmental & Resource Economists* 31:133–157.

Fowlie, Meredith. 2010. "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement". *American Economic Review* 100 (3): 837–869.

Frey, Elaine F. 2013. "Technology Diffusion and Environmental Regulation: The Adoption of Scrubbers by Coal-Fired Power Plants". *The Energy Journal* 34 (1): 177–205.

Karp, Larry S., and Christian P. Traeger. 2013. "Dynamic Methods in Environmental and Resource Economics". Unpublished.

Rampini, Adriano A., Amir Sufi, and S. Viswanathan. 2014. "Dynamic risk management". *Journal of Financial Economics* 111:271–296.

Cicala, Steve. 2015. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation". *American Economic Review* 105 (1): 411–444.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty". *The Quarterly Journal of Economics* 131 (4): 1593–1636.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant. 2018. "Hotelling under Pressure". *Journal of Political Economy* 126 (3): 984–1026.

Barron, Alexander R., et al. 2018. "Policy Insights from the EMF 32 Study on U.S. Carbon Tax Scenarios". *Climate Change Economics* 9 (1): 1840003-1 - 1840003–47.

McFarland, James R., et al. 2018. "Overview of the EMF 32 Study on U.S. Carbon Tax Scenarios". *Climate Change Economics* 9 (1): 1840002-1 - 1840002–37.

Ramstein, Céline, et al. 2018. *State and Trends of Carbon Pricing*. Washington, DC: World Bank.