DISCUSSION PAPER

April 2010; Revised June 2011 ■ RFF DP 10-27-REV

Soft and Hard Price Collars in a Cap-and-Trade System

A Comparative Analysis

Harrison Fell, Dallas Burtraw, Richard Morgenstern, Karen Palmer, and Louis Preonas

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

Soft and Hard Price Collars in a Cap-and-Trade System: A Comparative Analysis

Harrison Fell, Dallas Burtraw, Richard Morgenstern, Karen Palmer, and Louis Preonas

Abstract

We use a stochastic dynamic framework to compare price collars (price ceilings and floors) in a cap-and-trade system with uncertainty in the level of baseline emissions and costs. We consider soft collars, which provide limited volume of additional emission allowances (a reserve) at the price ceiling, and hard collars, which provide an unlimited supply of additional allowances, thereby preventing allowance prices from exceeding the price ceiling. Conversely, allowances are removed from the market if prices fall to the floor. We find that increasing the size of the reserve strictly lowers expected net present values of compliance costs; however, there is a diminishing effect as the allowance reserve is expanded. Most of the expected cost savings are achieved with a modest reserve. Consequently, a rather limited soft price collar could provide considerable assurance about cost while preventing the possibility that emissions could spiral out of control.

Key Words: climate change, cap-and-trade, price collars, stochastic dynamic programming

JEL Classification Numbers: Q54, Q58, C61

Contents

Soft and Hard Price Collars in a Cap-and-Trade System: A Comparative Analysis

Harrison Fell, Dallas Burtraw, Richard Morgenstern, Karen Palmer, and Louis Preonas*

I. Introduction

Concerns about potentially extreme allowance price and compliance cost outcomes have hampered efforts to adopt a U.S. cap-and-trade policy to regulate emissions of carbon dioxide (CO_2) and other greenhouse gases (GHG). Hybrid policies, in particular price collars, which add allowance price floors and ceilings, have attracted attention as a way to constrain potential costs and price variability in a cap-and-trade system. Price collars also appear to have gained political traction, having been included in various forms in recently proposed GHG mitigation bills (e.g., H.R.2454 (Waxman-Markey), S.2879 (Cantwell-Collins) and S.1733 (Kerry-Boxer)).

Hybrid policies have been considered in the economics literature for many years. Early works by Roberts and Spence (1976) and Weitzman (1978) considered price floors and ceilings along with emission quantity constraints in static models with uncertain environmental benefits and costs. The issue of price ceilings was brought up again with respect policies designed to mitigate climate change in Kopp et al. (1997) and explored more thoroughly in McKibben and Wilcoxen (1997) and Pizer (2002) by looking at dynamic models that combined price ceilings with a quantity regulations. Philibert (2008) and Burtraw et al. (2010) consider dynamic simulation models with uncertainty about abatement costs that layered price collars (both a price floor and ceiling) onto quantity constraints, while Fell and Morgenstern (2009) consider price collars in a stochastic dynamic framework that also allowed for allowance banking and borrowing. Webster et al. (2010) also consider a stochastic dynamic framework, but compare a hybrid policy with a price ceiling to a policy that allocates allowances based on emission intensity targets. Grüll and Taschini (2011) provide an analytic exposition of price floors and price ceilings from the options pricing perspective. A common thread to all of this research is that they consider price controls with limitless price support and generally find that hybrid policies offer welfare gains over pure quantity policies.

* Fell is a Fellow, Burtraw, Morgenstern, and Palmer are Senior Fellows, and Preonas is a Research Assistant at Resources for the Future, Washington, DC. This research was supported by a grant from National Commission for Energy Policy and Mistra's Climate Policy Research Program (Clipore).

Price collars may constrain costs, but in doing so they expand the range of possible emission outcomes, which calls into question the environmental integrity of such systems. Specifically, if hitting a price ceiling were to introduce additional emissions allowances, there is concern that if prices hit the ceiling frequently, it would not be possible to meet the emission reduction goals of the system. As a compromise between the concerns of environmental integrity and cost uncertainty, Murray et al. (2009) propose an allowance reserve of limited size in which there is a pre-set maximum number of allowances that can be added to the system in any given period when the price ceiling is hit. We describe this design as a *soft* price ceiling because if the additional allowances in the reserve are not sufficient to satisfy demand, the allowance price can rise above the price ceiling. In this paper, we analyze a similar limited-reserve price instrument in the context of a price collar on emissions. A price collar is a two-sided price instrument that includes an upper bound price support in the form of a price ceiling and a lower bound price support in the form of a price floor. The price floor insures that if allowance prices are low, abatement costs do not fall to unacceptably low levels and instead some of those cost savings are realized as additional environmental benefits. As designated here, the soft collar plan differs from the *hard* price collar, in which the price ceiling and price floor is strictly enforced in each period no matter how many allowances must be added to or removed from the system.[1]

Prominent legislative proposals for cap and trade in the US have featured versions of a mixed system with a limited allowance reserve (soft collar) triggered at a price ceiling, coupled with an unlimited withdrawal of allowances (hard collar) triggered at a price floor.[2] The key parameters of the policy – the price triggers and size of the allowance reserve – have been a central topic in the policy dialogue. Although the qualitative environmental and cost-containing tradeoffs between the hard and soft collars seem quite obvious, little has been done in the way of formal comparisons of the two systems.

[1] In principle, one could also consider different limits on the quantity of allowances that would be withdrawn to support the price floor, although this feature has not been a major issue in public debates.

[2] Recent federal proposals including H.R. 2454 (Waxman-Markey), S.2879 (Cantwell-Collins), S. 1733 (Kerry-Boxer), and state programs including the ongoing northeast Regional Greenhouse Gas Initiative (RGGI) and California's trading program slated to commence in 2012 include an auction for an important fraction of allowance allocations, with the reserve price in the auction serving as a price floor for allowance prices in the program. The quantity that would be withheld from the market in some cases is limited but sufficiently large to be nonbinding. With the exception of S.2879, most of these proposals and programs have no price ceiling. H.R. 2454 limited the amount the price could increase each year, and S. 1733 had a limited quantity reserve that would be sold at or above a price ceiling trigger. In every case price increases also are expected to be dampened by the generous ability to tap offset markets. In the case of RGGI, the ability to use offsets is triggered by a price ceiling.

In this paper, we use analytical and numerical methods to address this gap. We begin with an analytical representation of the problem in a static setting. This formulation demonstrates an explicit tradeoff between the variance and range of emissions and costs along a continuum from no allowance reserve to an infinite reserve (equivalent to a hard collar). Moreover, we show diminishing marginal returns to the size of the reserve. That is, the first allowance (equivalent to a ton of emissions) that is added to the reserve leads to the greatest cost savings, and each additional ton leads to additional cost savings albeit at a declining rate. Conversely, the first allowance that would be removed from the system at the price floor imposes the greatest incremental cost and the change in cost declines with additional tons. We show this finding holds under a symmetric or asymmetric collar.

These results focus attention on two specific questions. Does this straightforward tradeoff exist in a dynamic uncertain context, i.e. do the marginal benefits for cost management continue to decline with the size of the reserve where, for example, allowance banking and borrowing may alter the use of the allowance reserve? And, what is the empirical relevance of this tradeoff, i.e. is there a distinct cut-off point(s) or is there a broad range for the size of the allowance reserve that yields comparable benefits? More precisely, can a small allowance reserve with limited potential environmental consequence provide most of the benefits with respect to cost management? The answer to this question is important; an affirmative answer would suggest that the worst fears of environmental advocates – that an unlimited number of allowances might flood the system – can be assuaged with a soft collar, while also providing assurance to industry that costs are unlikely to spiral out of control.

To examine these questions in a policy-relevant context we evaluate the two systems over a variety of design specifications in a dynamic stochastic model, solved numerically with parameters germane to U.S. climate policy. The framework allows for the opportunity to bank or borrow emissions allowances over time and uncertainty about future baseline emission levels in the absence of regulation and, thus, future abatement costs. We compare the net present value (NPV) of regulatory costs, emission outcomes, and allowance price variability for hard and soft collar systems. We consider a range of collar widths (price ceiling minus price floor) and soft collar designs under various allowance reserve sizes. Importantly, for each specification we set the collars such that all policy designs lead to the same level of *expected* cumulative emission outcomes. By doing so, we are able to get a more apples-to-apples comparison of the different approaches, at least in terms of environmental integrity.[3]

Our results in the dynamic simulations are consistent with intuition developed in the static analysis. For the parameter settings examined, moving a soft-collar policy more towards a hard collar leads to lower expected NPV of compliance costs and lowers allowance price variability, but these benefits have diminishing returns. Most of this economic benefit can be achieved with a soft price collar that has a limited allowance buyback and reserve. Conversely, even if the expected outcome is held constant we find that increasing the allowance buyback and reserve can lead to greater variability in the environmental outcome. But, the greater variability comes at an increasing rate, which means there is relatively little environmental consequence associated with the initial introduction of an allowance buyback and reserve. In other words, there is a broad range of designs and values for a soft price collar that harvests most of the economic advantages in terms of reduced expected costs and price variability, but with little effect on environmental outcomes. Notionally, the initially large and decreasing marginal benefits of a soft price collar contrast with the initially small and growing marginal cost in terms of environmental concerns, suggesting an opportunity for compromise that optimizes a multi-attribute payoff function for climate policy.

Some forms of price controls, similar to the collars we discuss here, have been implemented in a variety of commodity markets. Economic analysis of these systems has focused on how limited price controls or limited buffer stocks could be manipulated. For example, Salant (1983) investigates the possibility of speculative attacks finding that a buffer stock would have to be infinitely large to withstand attack. Stocking (2010) applies these ideas to an emissions allowance market and argues that even an infinite allowance reserve could be subject to attack. We do not look at such noncompetitive behavior, but note that the opportunities for speculative attacks appear less applicable to emissions markets as they are outlined in existing proposals and programs than in other commodity markets. First, emissions allowances are not a physical commodity and the supply could be directly changed by policy makers. Second, allowances have zero storage cost and the existence of an intertemporal allowance bank raises the cost of trying to manipulate market equilibrium. Moreover, those firms holding a

[3] As discussed in Pizer (2002), the marginal damages curve for a stock pollutant such as CO_2 is generally considered to be quite flat. Thus, by comparing different collar designs in a CO_2 emissions context and conditional on generating the same expected emissions we are essentially considering a framework that holds expected emission damages constant. If one was to conduct a similar analysis on a pollutant that has a more convex marginal damage function, one should explicitly consider the effect of the variance of emissions on expected damages.

positive bank would be negatively affected by efforts to artificially drive down the allowance price, making collusion difficult. Further, proposals such as S. 1733 would implement allowance reserves within an auction, so the sales of reserve allowances do not occur at a fixed price as considered in the literature but at a market clearing price that is likely to rise above the trigger that opens the reserve. Nonetheless, while the existing literature has identified the potential disadvantages of limited price support, we address an advantage of limited measures that has not been identified previously, and which may have relevance beyond the design of an emissions allowance market: the biggest "bang for the buck" occurs with the introduction of an initial allowance reserve and the cost stabilizing feature of these controls has diminishing marginal value as the reserve is expanded.

The remainder of the paper is organized as follows. In Section II, we present an analytical model and derive general results in a static framework. We then turn to dynamic stochastic modeling to test the generality of these results in a dynamic setting with uncertainty. In Section III, we describe the simulation analysis and model parameterization. Section IV describes the results of the analysis. We discuss our conclusions in Section V.

II. Model Setup

As shown in Montgomery (1972), in a static context, cap-and-trade implemented in a competitive market can lead to the efficient emissions outcome. Therefore, we represent the emissions outcome of the regulated sectors using a representative-firm framework. We represent the unregulated emissions of the cost-minimizing representative firm as $\bar{q}+\theta$, where $\bar{q}$ is the expected unregulated emissions and $\theta \in [\theta_{min}, \theta_{max}]$ is a random shock to emissions that is symmetrically distributed about zero (i.e, $-\theta_{min} = \theta_{max}$). Under cap-and-trade regulation, the firm's allocation of emissions allowances is y (an equivalent formulation would allow for an auction of allowances). Abatement for the case with no price controls is $a=\bar{q}+\theta-y$. In determining the cost of the regulation, we are interested only in the expected abatement costs since allowance reserve sales and purchases represent transfers between the firm and the regulator. The cost of the abatement is $C(a)$, where C is an increasing convex abatement cost function ($C'>0$, $C''\geq 0$).

We assume price controls are implemented such that the regulator agrees to buy a limited quantity of allowances at or below the price floor (P^f) and to sell a limited quantity of additional allowances at or above the price ceiling (P^c). With price controls implemented, the firm must decide how many allowances to sell back (y^f) at P^f or to buy (y^c) at P^c. The firm's maximization problem becomes:

$$\max_{y^f, y^c} -C\left(\bar{q}+\theta-\left(y-y^f+y^c\right)\right)+P^{f*}y^f-P^{c*}y^c \tag{1}$$

subject to:

$$\begin{gathered} 0 \le y^f \le y^f_{\max} \text{ and } 0 \le y^c \le y^c_{\max} \\ P^{f*} = \min\left(P^f, C'\left(\bar{q}+\theta+y^f-y\right)\right) \\ P^{c*} = \max\left(P^c, C'\left(\bar{q}+\theta-y-y^c\right)\right) \end{gathered} \tag{2}$$

where $y^f_{\max}$ and $y^c_{\max}$ represent the maximum amount of allowances the firm is allowed to sell back at the price floor and buy at the price ceiling, respectively. Note also the definitions of P^{f*} and P^{c*} allow marginal abatement costs (i.e., allowance prices) to fall below P^f and rise above P^c, respectively.[4] The first order conditions for the maximization imply:

$$\text{if } y^f > 0 \text{ then } C'\left(\bar{q}+\theta-\left(y-y^f\right)\right) \le P^{f*} \tag{3}$$

and

$$\text{if } y^c > 0 \text{ then } C'\left(\bar{q}+\theta-\left(y+y^c\right)\right) \ge P^{c*} \tag{4}$$

which state that the firm will buy or sell allowances during periods when the price controls bind. The marginal cost of abatement can be less than the relevant price control in (3) when the buyback quantity is exhausted $\left(y^f = y^f_{\max}\right)$ or exceed the price control in (4) when the reserve quantity is exhausted $\left(y^c = y^c_{\max}\right)$. Defining $h(P) = C'^{-1}(P)$, the choice of y^f and y^c over the entire span of θ is:

$$y^f = \begin{cases} y^f_{\max} \text{ if } \theta_{\min} \le \theta < \theta^f_1 \\ h\left(P^f\right)+y-\bar{q}-\theta \text{ if } \theta^f_1 \le \theta \le \theta^f_2 \\ 0 \text{ if } \theta > \theta^f_2 \end{cases} \tag{5}$$

$$y^c = \begin{cases} 0 \text{ if } \theta_{\min} \le \theta < \theta^c_1 \\ \bar{q}+\theta-y-h\left(P^c\right) \text{ if } \theta^c_1 \le \theta \le \theta^c_2 \\ y^c_{\max} \text{ if } \theta > \theta^c_2 \end{cases} \tag{6}$$

The indicated values of θ are the shocks to baseline emissions that result in allowances prices that trigger the price floor or ceiling:

---

[4] Prices may fall below P^f, e.g. if the quantity of allowances the firm would like to sell back to the regulator at P^f exceeds $y^f_{\max}$. One way the regulator could solve the rationing problem would be to implement a reverse auction with a reserve price equal to the price floor. In this situation the auction clearing price could fall below the price floor. Similarly, to distribute allowances from the reserve the regulator could implement an auction with reserve price equal to the price ceiling.

$$
\begin{aligned}
\theta_1^f &= h\left(P^f\right) + y - y_{\max}^f - \bar{q} \\
\theta_2^f &= h\left(P^f\right) + y - \bar{q} \\
\theta_1^c &= h\left(P^c\right) + y - \bar{q} \\
\theta_2^c &= h\left(P^c\right) + y + y_{\max}^c - \bar{q}
\end{aligned} \tag{7}
$$

Given (5) and (6) and assuming a probability density function for θ of $f(\theta)$, the four critical values of θ in (7) determine five terms in the expected cost of abatement:

$$
\begin{aligned}
E\left[C(a)\right] = &\int_{\theta_{\min}}^{\theta_1^f} C(\bar{q} + \theta + y_{\max}^f - y) f(\theta) d\theta + \int_{\theta_1^f}^{\theta_2^f} C(h(P^f)) f(\theta) d\theta + \\
&\int_{\theta_2^f}^{\theta_1^c} C(\bar{q} + \theta - y) f(\theta) d\theta + \int_{\theta_1^c}^{\theta_2^c} C(h(P^c)) f(\theta) d\theta + \int_{\theta_2^c}^{\theta_{\max}} C(\bar{q} + \theta - y_{\max}^c - y) f(\theta) d\theta
\end{aligned} \tag{8}
$$

The elements of the expected cost function with a price collar are illustrated in Figure 1.[5] The top panel illustrates the range of emissions that occur given the realized value of θ. Under regulation, when $\theta = 0$ the expected price of emissions (marginal abatement cost) is $E[P]$. Allowance prices increase with higher values of θ (higher emissions). To illustrate the price ceiling, if $\theta \geq \theta_1^c$ the allowance price would trigger the price collar and additional allowances would be introduced to the market. For realizations of $\theta \in \left[\theta_1^c, \theta_2^c\right]$ the price equals the price collar (P^c), and for values of $\theta \geq \theta_2^c$ the price rises above P^c. The second panel transposes the level of abatement as a function of the changes to baseline emissions, and illustrates that as the size of the allowance buyback or reserve increases, the range of potential abatement outcomes is reduced. The third panel illustrates that for a static convex expected cost function a change in the range of potential abatement outcomes will change expected costs. The following lemma establishes a useful result.

Lemma 1: An increase in the quantity of the buyback at the allowance price floor increases expected abatement cost, and an increase in the quantity of the reserve at the allowance price ceiling reduces expected abatement cost.

[5] Note that (8) and the bottom panel of Figure 1 include abatement cost only and exclude the firm's cost reductions from sales at or below the price floor and the cost additions from purchases at or above the price ceilings. These additional benefits and cost to the firm represent transfers between the firm and regulator. As noted, since we focus on assessing the price collar design based on the abatement costs, we do not include these transfers in the cost assessments.

Corollary 1: If the collar is designed such $y^f_{\max} = y^c_{\max} = y_{\max}$ and the price levels are set such that the *ex ante* expected value E[$y + y^c - y^f$] = y, then an increase in $y_{\max}$ reduces expected abatement cost.

Proof: See appendix.

The next lemma establishes diminishing returns to scale in both the allowance buyback $\left(y^f_{\max}\right)$ and reserve $\left(y^c_{\max}\right)$ as long as the probability density of θ is single-peaked between the price floor and price ceiling.[6] That is, the reduction in expected costs from adding additional allowances to the reserve is diminishing with the scale of the reserve.

Lemma 2: The increase in cost from increasing the size of the buyback of allowances at the price floor declines with the scale of buyback, and the decrease in cost from increasing the reserve size declines as the reserve size gets bigger.

Corollary 2: If the collar is designed such that $y^f_{\max} = y^c_{\max} = y_{\max}$ and the price levels are set such that E[$y + y^c - y^f$] = y, then the decrease in cost from increasing $y_{\max}$ declines as $y_{\max}$ gets bigger.

Proof: See appendix.

The relevance of Lemma 2 is to show that as the size of the allowance reserve grows, there is a diminishing return from cost mitigation that is commensurate with an increasing risk of high emissions outcomes. In the policy dialogue, efforts to constrain mitigation costs have focused primarily on limiting the possibility of cost spikes by introducing a price ceiling and an allowance reserve, while environmental concerns have focused on the possibility that if the reserve is unbounded, i.e. a hard collar or safety valve, there would be a positive probability that emissions could far exceed the intended emissions cap. Lemma 2, along with Corollary 2, suggests the possibility of a political compromise. However, the results in a static context ignore important features of a dynamic stochastic problem.

Dynamic Model

The static model is useful in building intuition; however, proposed regulations outline emission goals over many years and often allow for some degree of intertemporal allowance

[6] Letting θ_m be the mode of θ, then $\theta^f_2 \le \theta_m < \theta^c_1$.

trading (i.e., banking and borrowing). Thus, evaluating the effectiveness of different price control regimes is inherently a dynamic optimization problem.

A key difference between the static model and a dynamic model that allows banking and borrowing is that emissions in any given period, q_t, will no longer simply be $y_t - y_t^f + y_t^c$. The stochastic dynamic optimization problem for the representative firm can be written as:

$$\max_{q,y^c,y^f} \sum_{t=1}^{T} \beta^{t-1} E_t \left[-C_t(\bar{q}_t + \theta_t - q_t) + P_t^{f*} y_t^f - P_t^{c*} y_t^c \right] \tag{9}$$

subject to:

$$B_{t+1} = B_t + y_t - y_t^f + y_t^c - q_t \tag{10}$$

$$B_t \geq B_{\min,t} \tag{11}$$

$$\theta_t = \rho\theta_t + \varepsilon_t \tag{12}$$

and the boundary conditions for the allowance buyback and reserve, along with the price collar definitions, given in (2),where β is the discount factor, E_t is the expectations operator at time t, B_t is the bank value at t, and ε_t is an iid random term.[7] Constraint (10) states that the bank at the beginning of the next period is equal to the bank at the beginning of the current period, adjusted by the intertemporal trading ratio, plus allocations, less allowances sold at the floor price, plus allowances purchased at the ceiling less emissions. Equation (11) sets a constraint on the amount of borrowing, where $B_{\min,t}$, the amount of allowance debt that possible in period t, is defined by the regulator (e.g., $B_{\min,t} \leq 0$). Finally, condition (12) allows for shocks to baseline emissions to be correlated over time.

The Bellman equation for solving a dynamic optimization problem in discrete time is given as:

$$V\left(B_t, \theta_t\right) = \max_{q_t, y_t^c, y_t^f} \begin{pmatrix} -C_t(\bar{q}_t + \theta_t - q_t) - P_t^{c*} y_t^c + P_t^{f*} y_t^f \\ +\beta E_t \left[V_{t+1}\left(B_{t+1}, \theta_{t+1} \mid B_t, \theta_t, q_t, y_t^c, y_t^f\right)\right] \end{pmatrix} \tag{13}$$

The first order conditions for optimality are:

---

[7] Rubin (1996) extended Montgomery (1972) to show that with emissions trading and allowance banking and borrowing, a competitive permit market outcome will be the same as the lowest-cost outcome, allowing us to model the market as a single cost-minimizing firm. Additionally, we are assuming allowances are intertemporally traded on a one-for-one basis. See Kling and Rubin (1997) and Leiby and Rubin (2001) for discussions on the use of intertemporal trading ratios. Note also that we do not model endogenous technological change as this is beyond the scope of the current analysis. See Weber and Neuhoff (2010) for a two-period model of investment within a hybrid instrument context with hard price constraints.

$$C_t'\left(\bar{q}+\theta_t-q_t\right)=\beta E_t\left[\frac{\partial V_{t+1}(\bullet)}{\partial B_{t+1}}\right] \tag{14}$$

$$P_t^{f*}-\mu_1^f+\mu_2^f=\beta E_t\left[\frac{\partial V_{t+1}(\bullet)}{\partial B_{t+1}}\right] \tag{15}$$

$$P_t^{c*}-\mu_1^c+\mu_2^c=\beta E_t\left[\frac{\partial V_{t+1}(\bullet)}{\partial B_{t+1}}\right] \tag{16}$$

$$\frac{\partial V_t(\bullet)}{\partial B_t}=\beta E_t\left[\frac{\partial V_{t+1}(\bullet)}{\partial B_{t+1}}\right]+\lambda_t \tag{17}$$

where μ_j^i ($i = f, c$ and $j = 1, 2$) correspond to the Lagrange multipliers associated with the quantity constraints on y^f and y^c (refer to Figure 1); and λ_t is the Lagrange multiplier on the borrowing constraint. Given (14) and (17) we can formulate an Euler equation for the transition between periods as:

$$C_t'\left(\bar{q}+\theta_t-q_t\right)=E_t\left[\beta C_{t+1}'\left(\bar{q}_{t+1}+\theta_{t+1}-q_{t+1}\right)+\lambda_{t+1}\right] \tag{18}$$

The Euler equation gives a modified Hotelling rule, similar to that given in (J. Rubin, 1996). More specifically, we see from (18) that, in the absence of borrowing restrictions and price controls, the expected marginal cost of abatement (i.e. expected allowance prices) will rise at the discount rate as the Hotelling rule predicts. However, the inclusion of a borrowing restriction and/or price controls can make expected allowance prices diverge from the Hotelling path. For instance, if the borrowing restriction is binding the solution will depart from the Hotelling path. Furthermore, if the borrowing limit has a positive probability of being reached in the next period then $E_t[\lambda_{t+1}] > 0$, current emissions must decline to preserve (18). Thus, the existence of a borrowing constraint makes the solution to the problem diverge from a certainty equivalence form, i.e. the solution evaluated at the mean of the uncertain parameters, as well as provides an incentive for a form of precautionary savings even if the firm is risk neutral (assuming shocks are such that $E_t[\lambda_{t+1}] > 0$ for some t) .[8]

From (15), we see that whenever additional allowances are being purchased from the reserve such that the reserve limit is not met ($0 < y_t^c < y_{\max}^c$ so that $\mu_1^c = \mu_2^c = 0$), y_t^c will be set such that next period's expected marginal value of a banked allowance equates to the ceiling

[8] The notion that borrowing constraints may create a form of precautionary savings is mentioned briefly with regards to emissions regulation in Schennach (2000). More generally, the role of ad-hoc borrowing constraints such as the one modeled here have been considered at length in macroeconomic consumption and savings models (e.g. Zeldes, 1989; Deaton, 1991; and Aiyagari, 1994). Note also, our model formulation implies incomplete markets, forcing the firm to self-insure.

price, P_t^c. From (14) we also see that marginal costs of abatement will be set equal to P_t^c during periods when $0 < y_t^c < y_{\max}^c$, which again will cause a divergence between the standard Hotelling path and possible price paths obtained in this problem. As with the static case when $\theta_1^c \leq \theta \leq \theta_2^c$, for a given bounded θ_t with values possible to make the price ceiling binding we see an increase in the allowance reserve will reduce the upper bound of abatement in a given period because abatement is constant over the range of θ_t that makes $0 < y_t^c < y_{\max}^c$ (analogous to Figure 1). Likewise, based on (16) and (14), an increase in the buyback will increase the lower bound of abatement.

Unfortunately, this problem does not allow one to solve for fully analytic policy functions and, thus, categorizations of the abatement costs. To see what effect reserve size, buyback limits, and other regulation design parameters have on abatement costs, emission outcomes, and other model outcomes, we parameterize the model and discretize the state space so that we can solve it numerically. Such an approach limits the generality of our results, though as we discuss below, many of our results are in line with the generalized static model presented above. Furthermore, by parameterizing our numeric model with values relevant to potential U.S. CO_2 regulation or legislation, we are able to provide policy relevant findings.

II. Numerical Model

Our numerical method is to solve the problem through backward recursion. To begin, however, we must first impose a quadratic functional form for $C(a_t)$:

$$C_t\left(\bar{q}_t + \theta_t - q_t\right) = \frac{c_t}{2}\left(\bar{q}_t + \theta_t - q_t\right)^2 \tag{19}$$

The quadratic form of (19) is similar to that used in (Richard Newell et al., 2005), except that it includes an additional squared θ_t term. We must also impose a terminal condition. In this case, we force all allowance debt to be repaid in the last period, so $B_{T+1} \geq 0$, with $B_{T+1} = 0$ for all cases where $B_T + y_T < \bar{q}_T + \theta_T$.

To discretize the state space of the model, the state variables B_t and θ_t are allowed to take on N_B and N_θ possible values, respectively.[9] Based on equation (12), we can form a probability transition matrix to describe the probability of moving from any discrete θ_t value to any other

[9] All other time-dependent state variables (c_t and $\bar{q}_t$) are assumed to evolve in a deterministic manner exogenous to the control variables q_t, y_t^c, and y_t^f.

θ_{t+1} value.[10] Given the terminal condition for B_{T+1} and bank transition equation (10), we can solve for $N_B \times N_\theta$ matrices of q_T, y_T^f, and y_T^c for all given state outcomes (B_T, θ_T) as the values that maximize $V_T(B_T, \theta_T) = -C_T(\bar{q}_T + \theta_T - q_T) + P_t^f y_T^f - P_T^c y_T^c$. Knowing $V_T(B_T, \theta_T)$ for all possible states and the probability transition matrix for θ, we can then determine $E_{T-1}\left(V_T\left(B_T, \theta_T\right) | B_{T-1}, \theta_{T-1}\right)$. The values $q_{T-1}(B_{T-1}, \theta_{T-1})$, $y_{T-1}^f(B_{T-1}, \theta_{T-1})$, $y_{T-1}^c(B_{T-1}, \theta_{T-1})$ and consequently $V_{T-1}\left(B_{T-1}, \theta_{T-1}\right)$, can be solved as the values that maximize (13). We continue this backward iteration process of solving for each period the $q_{t-1}(B_{t-1}, \theta_{t-1})$, $y_{t-1}^f(B_{t-1}, \theta_{t-1})$, and $y_{t-1}^c(B_{t-1}, \theta_{t-1})$ that maximize (13) given $E_{t-1}\left(V_t\left(B_t, \theta_t\right) | B_{t-1}, \theta_{t-1}\right)$ until we iterate back to the initial period.

The end result of this backward recursion is a set of T matrices with dimension $N_B \times N_\theta$ for q_t, y_t^c, and y_t^f that solve the optimal quantities for these variables for any given state realization in any given time period (i.e., the optimal policy rules for value maximization, i.e. cost minimization). Knowing the optimal sequence for the control variables and the bank state dynamics given in (10), we can conduct simulation analyses. We first generate N_{sim} different simulation paths for the emissions shock θ, where a path is a 1 x T vector of outcomes based on the transition equation (12). Given the simulated shock paths and an initial bank condition, $B_1 = 0$, we can use the optimal control matrices of q_t, y_t^c, and y_t^f to map out N_{sim} emission, offset, and bank path realizations. The path realizations become the basis for the cost, cumulative emission, cumulative offsets, and price variability metrics we use to compare various policy designs under a range of parameter specifications.[11]

Model Parameterization

The model is keyed primarily to H.R. 2454 and to the analysis of this bill performed by the U.S. Department of Energy's Energy Information Administration EIA (2009).

To begin, we set the terminal period of the model at $T = 39$ to simulate H.R. 2454's 2012 – 2050 timeframe. The emission allowances, y_t, are the summation of the allocations to various entities taken directly from H.R. 2454 as laid out in Section 702 of the bill plus expected offset

[10] The discretized B_t is evenly spaced between user-defined $B_{\min,t}$ and $B_{\max,t}$ values. The discretization step and probability transition matrix formulation for θ is done using the discretization method for AR(1) processes described in Adda and Cooper (2003).

[11] The program used to generate the optimal policy rules and run the simulations, along with all other variables used within the simulations, are available from the authors upon request.

purchases as given in EIA (2009).[12] Borrowing is assumed to be prohibited ($B_{\min,t} = 0$). No limit is given for the maximum bank level, but for discretization we set it at 45 giga-tons of CO_2 equivalent ($GtCO_2e$), which was sufficiently high to not be binding for all simulations conducted. For the expected baseline emissions path, $\bar{q}_t$, we use the baseline emissions path given in EIA (2009) for 2012 – 2030, EIA's analysis period, and extend the trend of EIA's baseline emissions from 2025 – 2030 for the remaining years of our analysis. As stated above, we assume an exogenous decline in the slope of the marginal abatement cost curve. Unfortunately, to our knowledge, there is no direct corollary to this decline rate that has been empirically estimated or estimated via simulation. We set the decline rate, g_c, at what we feel is a modest -1.25% annual rate. We set the initial value of c_0 = \$62/$mtCO_2e$ per $GmtCO_2e$ to approximate the price path in EIA (2009) for EIA's low discount rate case, which sets the discount rate to 0.05, with no uncertainty in the model ($\theta_t = 0, \ \forall \ t$).

For the price control systems, we consider three general design specifications, the symmetric soft collar, the one-sided soft collar, and the hard collar. In our symmetric soft collar we consider price controls where the buyback limit at the price floor is equal to the reserve size at the price ceiling. That is, the maximum quantity of allowances the regulator has agreed to purchase at P_t^f is equal to the maximum quantity of reserve allowances it will sell at P_t^c. The one-sided soft collar design is set such that the buy-back quantity is unlimited (i.e., there is a hard price floor) but there is a finite reserve of allowances available for purchase at the price ceiling. As described above, for the hard collar systems, the regulator has a standing offer to buy an unlimited quantity of allowances at P_t^f and offers an unlimited quantity of allowances for sale at P_t^c. In all specifications, we design the collars such that both the floor and the ceiling rise at the discount rate, $P_t^i = P_1^i(1+r)^{t-1}$ for $i = f, c$. We also take the width of the initial collar, $X_1 = P_1^c - P_1^f$, as a pre-determined variable and consider X_1 = (10, 15, 20). As mentioned above, initial values of the collars, P_1^f and P_1^c, are calculated through an iterative process such that the *expected* cumulative emissions, $\sum_{t=1}^{T} E[q_t]$, are equal to the cumulative allocation, $\sum_{t=1}^{T} y_t$ which is constant across scenarios. For the symmetric soft collar case we set $y_{\max,t}^f = \alpha y_t$ and for both soft

---

[12] The EIA (2009) analysis considers several different offset provision scenarios. We present results using the offset provisions from EIA's (2009) "High Offset" scenario. This scenario gives annual offsets provisions for 2012 – 2030. We extrapolate the remaining years offset provisions based on the trend of the last five years of offsets provisions presented in EIA (2009). We also ran models with the EIA's "Base" and "Low International Offset" scenarios. The results are qualitatively the same as presented here and can be made available from the authors upon request.

collar cases we set $y^c_{\max,t} = \alpha y_t$, where α = 0.05, 0.10, 0.20, 0.30 and 0.40 in our core scenarios. These parameters, as well as all other relevant parameters, are presented in Table 1.

IV. Results

A summary of the expected NPV of abatement costs for the different collar designs is given in Table 2.[13] The first column of the table gives the initial width of the collar, while the remaining columns give the expected NPV of abatement costs, with 95 percent intervals given in brackets below, for various α settings (i.e., reserve sizes) given as the header of each column.[14] Note that $\alpha = \infty$ corresponds to the hard collar case and the results with no price controls at all are given at the top of the table (no collar case). As expected, the table shows that with an increasing reserve size, expected abatement costs decline. We also see a decline in the range of possible cost outcomes, as shown by the narrowing of the 95 percent intervals.

The rate of decline in abatement cost is clearly represented in Figure 2. Here we plot the share of potential expected abatement cost savings achieved by moving from a no-collar case to a hard collar. Figure 2 shows that for each of the collar width settings and across both of the soft-collar types, increasing the reserve size yields diminishing abatement cost-savings. Furthermore, for a given type of soft collar, symmetric or non-symmetric, the rate of decline in abatement cost is quite similar across the different collar widths. However, the differences in the rates of decline between the non-symmetric and the symmetric soft collars are considerable. The non-symmetric collar achieves almost half of the cost savings between the no-collar and hard-collar cases with a modest annual reserve size equal to five percent of the annual reserve; and approximately 80 percent of those cost savings are achieved with annual reserve sizes equal to 20 percent of the allocation. On the other hand, symmetric collars need annual reserves equal to 40 percent of the allocation to be able to achieve 80 percent of the cost difference between no-collars and a hard-collar.

What drives the differences between these abatement cost-saving returns with increasing reserve sizes for the symmetric and non-symmetric collars? Recall that we have imposed the

[13] Note that the cost values given in Table 2 include only abatement costs and neglect costs the firm experience in buying additional allowances at the price ceiling or proceeds it gets from selling allowances at the price floor. Again, these additional costs and benefits, while accounted for in the optimization routine, represent transfers between the regulator and firm, and, thus, do not represent a true additional cost or benefit of the system.

[14] In our discussion of the numeric results, we generically refer to “reserve” size to simultaneously mean the reserve limit and buyback limit for the case of the symmetric collar because $y^f_{\max,t} = y^c_{\max,t}$.

expected environmental equality condition that forces all collar designs to yield the same expected cumulative emissions. Recall also that the non-symmetric collar cases do not limit the quantity the firm can sell to the regulator at the price floor and from (14) and (15) we know that regardless of what the reserve size is in the non-symmetric case the firm will always adjust actual emissions to equate marginal abatement costs to the price floor when the price floor is triggered. Therefore, in order to maintain our environmental equality condition, the set of price collars for a given reserve size $\left(P_1^i(\alpha),...,P_T^i(\alpha);\ i=f,c\right)$ must strictly be declining as the reserve size is reduced ($P_t^i(\alpha_1)<P_t^i(\alpha_2)\forall i,t$ if $\alpha_1<\alpha_2$), assuming the shocks to baseline emissions are bounded such that both the floor and the ceiling are met with positive probabilities. This result can be seen in Table 3, which shows the initial price ceiling for each reserve size parameter, α, and initial width.[15] Therefore, beyond the cost convexity that creates diminishing cost saving returns in the static case, the rate at which the cost savings diminish is accelerated (i.e., the second derivative becomes more negative) because the cost collars are shifting up as the reserve size is increasing.

On the other hand, note from Table that the initial price ceiling levels for the symmetric type soft collars do not strictly increase as the reserve size increases and are relatively stable across the various reserve sizes given. One should expect such a result, given that the shocks are evenly distributed about zero and increasing the reserve size equally increases the limit of additional allowances that can enter the system from reserve purchases and the limit of allowances that can be taken out of the system from the regulator buyback. Indeed, the only reason that we don't see the exact same collar set for various reserve sizes in the symmetric collar case is due to the complex way in which reserve sizes alter future cost and banking constraint expectations. The altered expectations give rise to altered emissions and additional allowance sales/purchases at the price floors/ceilings as described in equations (14) – (16) and (18). Given that for the symmetric-collar cases the collar sets remain relatively constant for the various α's, the diminishing returns to increasing the reserve size will be primarily driven by the convex nature of the cost function and not by accelerated increasing collar sets as seen with the non-symmetric collars.

With respect to environmental damages, marginal damages from stock pollutants such as CO_2 are often considered to be relatively flat, making the mean of the range of emissions the primary driver of environmental damages. However, considering the range of emission outcomes

---

[15] The initial price floors can be found by subtracting the initial width from the initial price ceiling. Note that since we assume that both the price floor and ceiling rise at the discount rate r for all collar settings, if $P_1^c(\alpha_1)<P_1^c(\alpha_2)$ then the price collar set associated with α_1 will always be below that of α_2.

is important for political reasons and potentially from an environmental damages perspective if one considers damage functions with large threshold effects. We, therefore, present 95 percent intervals of the cumulative emissions for both the symmetric and non-symmetric collar cases over an increasing range of reserve sizes and for the three different initial collar widths considered.[16] Across both symmetric and non-symmetric collar types, we see that increasing the reserve size increases the range of emission outcomes, though this increase levels off as the reserve size grows. One would expect such a result given the bounded nature of the baseline emissions shock. We also see for both collar types that increasing the width of the collar reduces the range of the emissions; again, as expected given that wider collars lead to less frequently accessed price triggers and thus fewer additional allowances entering the market and fewer allowances sold back to the regulator.

In comparing emission ranges across the non-symmetric and symmetric collar types, we see that while the upper bounds of emissions for the varying reserve sizes are roughly equivalent across collar types, the lower bounds of emissions differ considerably (see Figure 3). The non-symmetric collars lead to a much smaller lower bound of emissions than their symmetric counterparts, particularly at small reserve sizes. To clarify this observation, we present histograms of emission outcomes from the case where $\alpha = 0.05$ and $X = \$10/tCO_2$ under both the non-symmetric collar and symmetric collar in Figure 4. Note that for the non-symmetric case, we get a non-symmetric distribution with considerable bunching near the upper bound of emissions (total allocation plus total allowance reserve) and a long tail toward the lower end of emissions. This skewed distribution is due to the non-symmetric collar design, with its limitless allowance buyback and lower price collars required to meet our environmental equivalence restriction. Conversely, the symmetric soft collar histogram reveals a much more symmetric distribution about the mean emission level. However, there is still some bunching near the upper bound of emissions. Thus, with limited upper bound reserves, both the symmetric and non-symmetric collar types will have upper end emissions over a 95 percent interval near the upper-bound of allowable emissions. But, the limitless buyback of the non-symmetric collars, and the resulting lower collar sets, will force a skewed possible cumulative emissions distribution with a long left tail, resulting in a much smaller lower-bound over the 95 percent interval.

---

[16] We exclude the no collar cases from the figure. The no-collar case had a mean cumulative emissions level of 202.7 $GtCO_2$ with a 95 percent interval range from 195.8 – 203.2 $GtCO_2$. Note that mean of cumulative emissions from the no-collar case is less than the cumulative allocation of 203.2 $GtCO_2$. This finding occurs because with banking allowed in a model with uncertainty it is possible for the firm to have accumulated a bank that it cannot drive to zero by period *T*.

Finally, since a primary impetus for the use of price collars allowance is to create more stable allowance price, we also report a measure of allowance price variability for different collar specification in Table 4. Allowance price variability is captured by the root mean squared error of the price growth rate (RMSE):

$$RMSE = \frac{1}{N_{sims}} \sum_{i=1}^{N_{sims}} \sqrt{\frac{1}{T-1} \sum_{t=2}^{T} (P_{it}^{g} - \bar{P}_{t}^{g})^2}$$

where $P_{it}^{g} = \frac{(P_{it} - P_{it-1})}{P_{it-1}} \times 100$, $\bar{P}_{t}^{g} = \frac{1}{N_{sims}} \sum_{i=1}^{N_{sims}} \frac{(P_{it} - P_{it-1})}{P_{it-1}} \times 100$ and P_{it} is set to the marginal cost of abatement for simulation *i* in period *t*.[17] In general, we find that the RMSE values decrease with an increasing reserve, which is as expected since larger reserves reduce the range of possible emission prices (see Figure 1, panel 1). We also see that for both collar types, volatility decreases for a given reserve size as the collars become tighter (*X* gets smaller). Again this decrease is as expected because tighter collars will be triggered more frequently and, thus, limit variability.

In comparing variability measures across collar types, we see the RMSE of price growth under the non-symmetric soft collars more rapidly approach that of the hard collars as α increases. This finding is again due to the effect that our environment equality restriction has on the collar levels under e two collar types. As discussed above, the non-symmetric collars must be set at higher levels as α increases to maintain the constant expected cumulative emissions. The price ceilings will be triggered more frequently for the non-symmetric collars at lower reserve sizes because of these lower collars. Thus, while a larger reserve reduces price volatility by constraining possible allowance prices, this price-volatility-reducing advantage is partially offset for the non-symmetric collars because the smaller reserve will have lower, more accessible, price ceilings. Conversely, the collar levels of the symmetric-type collars are more consistent across reserve sizes. The price-volatility-reducing advantages of a larger reserve are, therefore, not offset with lower price collar levels under a system with a smaller reserve. Therefore, for the symmetric collars we see less price volatility with a larger reserve.

[17] We use variability in price growth rates as opposed to variability in levels to account for differences in the scale of marginal abatement costs that exist under different collar specifications.

V. Conclusion

This paper reviews the performance of hybrid emission control regulations that provide varying degrees of support for emission allowance prices. We begin by looking at so-called soft collars, hybrid policies with limited support of price ceilings and price floors. In a static model with a general convex abatement cost function and a symmetrically distributed shock in the level of baseline emissions, we find that increasing the support of the price ceiling (i.e., increasing the reserve limit) lowers expected abatement cost, but at a diminishing rate. Similarly, we find that increasing support for the price floor (i.e., increasing the buyback limit) increases expected abatement costs at a diminishing rate. Simultaneously and equally increasing the reserve and buyback limits decreases expected costs, but at a diminishing rate for a system where the collars are set such that expected emissions equal the allocation of allowances.

While these results are informative, actual systems contain more complex dynamics. We therefore investigate the performance of a soft collar in a stochastic dynamic framework that permits intertemporal allowance trading and correlated uncertainties. A full-fledged analytic solution to this problem is not possible, so we solve it numerically with parameters keyed to values relevant to the U.S. climate policy debate. We explore two primary collar types. One is a symmetric soft collar, which in any given period allows a limited sale of additional allowances into the system and an equally limited quantity which the regulator can buyback. The second is a non-symmetric soft collar, which allows a limited sale of additional allowances but unlimited buyback purchases. The two collar types are explored under a variety of reserve and buyback limits and collar widths. In all our calculations, the levels of the price collars are set such that expected cumulative emissions equal the cumulative allocation, so that environmental performance is equivalent on an expected value basis.

Our dynamic analysis confirms intuition built on the analytical results in the static model. As we increase the size of the buyback and reserve we move from a soft-collar setting with limited price support toward a hard collar setting with unlimited (strict) price support. In doing so, we find the range of possible emission outcomes increases at an increasing rate. This result illustrates the concern of environmental advocates who fear that a hard price collar could lead to outcomes with emissions in excess of the program goal. At the same time, business interests are concerned about the variability of emissions allowance prices, as well as the level of overall costs. We find that moving toward a hard collar reduces the variability in allowance prices, but at a diminishing rate. Most importantly, we also find that moving toward a hard collar decreases expected abatement costs, again at a diminishing rate, just as in the static analysis. That is, most of the cost savings that can be achieved with strict price controls can be achieved with rather modest reserve and buyback limits, particularly for the non-symmetric collar case.

These analytical and simulation results suggest that soft collars present an opportunity for compromise between stakeholders concerned about price volatility and system cost and those concerned about environmental outcomes. Most of the price and cost-related benefits that can be achieved under a hard price collar are achieved under a soft collar; and, the soft collar limits the range of environmental outcomes that might obtain under a hard price collar.

Appendix

From (8) we have the expected cost function:

$$E\left[C\left(a\right)\right]=\int_{\theta_{\min}}^{\theta_1^f} C(\bar{q}+\theta+y_{\max}^f-y)\,f\left(\theta\right)d\theta+\int_{\theta_1^f}^{\theta_2^f} C(h(P^f))\,f\left(\theta\right)d\theta+$$
$$\int_{\theta_2^f}^{\theta_1^c} C(\bar{q}+\theta-y)\,f\left(\theta\right)d\theta+\int_{\theta_1^c}^{\theta_2^c} C(h(P^c))\,f\left(\theta\right)d\theta+\int_{\theta_2^c}^{\theta_{\max}} C(\bar{q}+\theta-y_{\max}^c-y)\,f\left(\theta\right)d\theta \tag{A1}$$

We can now determine the effect of an increase in the buy-back limit ($y_{\max}^f$) and/or the increase in the allowance reserve ($y_{\max}^c$). We assume $\frac{\partial\theta_1^f}{dy_{\max}^f}<0$ and $\frac{\partial\theta_2^f}{dy_{\max}^f}=0$, and letting θ_m be the mode of θ, then $\theta_2^f \le \theta_m < \theta_1^c$. For an increase in $y_{\max}^f$,

$$\begin{aligned}\frac{\partial E\left[C\left(a\right)\right]}{\partial y_{\max}^f}&=C(\bar{q}+\theta_1^f+y_{\max}^f-y)\,f(\theta_1^f)\frac{\partial\theta_1^f}{\partial y_{\max}^f}+\int_{\theta_{\min}}^{\theta_1^f} C'(\bar{q}+\theta+y_{\max}^f-y)\,f\left(\theta\right)d\theta\\&\quad -C(h(P^f))\,f(\theta_1^f)\frac{\partial\theta_1^f}{\partial y_{\max}^f}\\&=\int_{\theta_{\min}}^{\theta_1^f} C'(\bar{q}+\theta+y_{\max}^f-y)\,f\left(\theta\right)d\theta\ge 0\end{aligned} \tag{A2}$$

and

$$\begin{aligned}\frac{\partial^2 E\left[C\left(a\right)\right]}{\left(\partial y_{\max}^f\right)^2}&=C'\left(\bar{q}+\theta_1^f+y_{\max}^f-y\right)f(\theta_1^f)\frac{\partial\theta_1^f}{\partial y_{\max}^f}+\int_{\theta_{\min}}^{\theta_1^f} C''(\bar{q}+\theta+y_{\max}^f-y)\,f\left(\theta\right)d\theta\\&=-C'\left(\bar{q}+\theta_1^f+y_{\max}^f-y\right)f(\theta_1^f)+C'\left(\bar{q}+\theta_1^f+y_{\max}^f-y\right)f(\theta_1^f)\\&\quad -C'\left(\bar{q}+\theta_{\min}+y_{\max}^f-y\right)f(\theta_{\min})-\int_{\theta_{\min}}^{\theta_1^f} C'(\bar{q}+\theta+y_{\max}^f-y)\,f'\left(\theta\right)d\theta\\&=-C'\left(\bar{q}+\theta_{\min}+y_{\max}^f-y\right)f(\theta_{\min})-\int_{\theta_{\min}}^{\theta_1^f} C'(\bar{q}+\theta+y_{\max}^f-y)\,f'\left(\theta\right)d\theta\le 0\end{aligned} \tag{A3}$$

From (A2) we see that as $y_{\max}^f$ increases, the expected cost of abatement increases (Lemma 1 for the price floor case). This result makes sense intuitively because a larger buy-back limit will allow the firm to sell more allowances back to the regulator and consequently abate more. From (A3), as long as the probability density function of θ is increasing or constant over the range of $\theta_{\min}$ to θ_1^f , (i.e. extreme values of θ are not more likely) the second derivative of expected costs with respect to $y_{\max}^f$ is negative. This relationship implies that the abatement cost gains from increasing $y_{\max}^f$ are falling as $y_{\max}^f$ continues to rise (Lemma 2 for the price floor case).

For an increase in the allowance reserve size, $y_{\max}^c$, the changes in expected costs are

$$\frac{\partial E\left[C(a)\right]}{\partial y_{\max}^c} = C(h(P^c))\, f(\theta_2^c)\frac{\partial \theta_2^c}{\partial y_{\max}^2} - C(\bar{q}+\theta_2^c - y_{\max}^c - y)\, f(\theta_2^c)\frac{\partial \theta_2^c}{\partial y_{\max}^c}$$

$$+ \int_{\theta_2^c}^{\theta_{\max}} -C'(\bar{q}+\theta - y_{\max}^c - y)\, f(\theta)d\theta \qquad \text{(A4)}$$

$$= -\int_{\theta_2^c}^{\theta_{\max}} C'(\bar{q}+\theta + y_{\max}^c - y)\, f(\theta)d\theta \le 0$$

and

$$\frac{\partial^2 E\left[C(a)\right]}{\left(\partial y_{\max}^c\right)^2} = C'\left(\bar{q}+\theta_2^c - y_{\max}^c - y\right) f(\theta_2^c)\frac{\partial \theta_2^c}{\partial y_{\max}^c} + \int_{\theta_2^c}^{\theta_{\max}} C''(\bar{q}+\theta - y_{\max}^c - y)\, f(\theta)d\theta$$

$$= C'\left(\bar{q}+\theta_2^c - y_{\max}^c - y\right) f(\theta_2^c) + C'\left(\bar{q}+\theta_{\max} + y_{\max}^c - y\right) f(\theta_{\max})$$

$$- C'\left(\bar{q}+\theta_2^c - y_{\max}^c - y\right) f(\theta_2^c) - \int_{\theta_2^c}^{\theta_{\max}} C'(\bar{q}+\theta - y_{\max}^c - y)\, f'(\theta)d\theta \qquad \text{(A5)}$$

$$= C'\left(\bar{q}+\theta_{\max} + y_{\max}^f - y\right) f(\theta_{\max}) - \int_{\theta_2^c}^{\theta_{\max}} C'(\bar{q}+\theta - y_{\max}^c - y)\, f'(\theta)d\theta$$

$$= C'\left(\bar{q}+\theta_2^c + y_{\max}^f - y\right) f(\theta_2^c) + \int_{\theta_2^c}^{\theta_{\max}} C''(\bar{q}+\theta - y_{\max}^c - y)\, f(\theta)d\theta \ge 0$$

Equation (A4) shows us that as the allowance reserve size increases, expected abatement costs decline (Lemma 1 for the price ceiling case). Again, this finding makes intuitive sense because as the reserve size increases firms will buy more allowances when baseline emission shocks are high and consequently abate less. Equation (A5) shows us the second derivative is positive implying that the cost decreases from increasing the reserve size are getting smaller as the reserve size gets bigger (Lemma 2 for the price ceiling case).

If the collars are designed symmetrically such that $y_{\max}^f = y_{\max}^c = y_{\max}$ and E[$y + y^c - y^f$] = y, then $-\theta_1^f = \theta_2^c$ and $-\theta_2^f = \theta_1^c$ because θ is distributed about zero. Given this design,

$$\frac{\partial E\left[C(a)\right]}{\partial y_{\max}} = -\int_{\theta_2^c}^{\theta_{\max}} C'(\bar{q}+\theta + y_{\max} - y)\, f(\theta)d\theta + \int_{\theta_{\min}}^{\theta_1^f} C'(\bar{q}+\theta + y_{\max} - y)\, f(\theta)d\theta \qquad \text{(A6)}$$

Because the marginal cost is increasing in abatement, $\frac{\partial E\left[C(a)\right]}{\partial y_{\max}} \le 0$. That is, (A6) shows us that for the case with a symmetric reserve and buyback limit and a collar designed to keep expected emissions at the allocation y, abatement costs decrease with a simultaneous and equal increase in the reserve and buyback limits (Corollary 1). We can also show that

$$\frac{\partial^2 E\left[C(a)\right]}{\left(\partial y_{\max}\right)^2} = C'\left(\bar{q} + \theta_{\max} + y^f_{\max} - y\right) f(\theta_{\max}) - \int_{\theta^c_2}^{\theta_{\max}} C'(\bar{q} + \theta - y^c_{\max} - y) f'(\theta) d\theta$$
$$- C'\left(\bar{q} + \theta_{\min} + y^f_{\max} - y\right) f(\theta_{\min}) - \int_{\theta_{\min}}^{\theta^f_1} C'(\bar{q} + \theta + y^f_{\max} - y) f'(\theta) d\theta \qquad \text{(A7)}$$

Given the assumed symmetry of the distribution of θ and the assumption that the distribution is increasing over the range $\theta_{\min}$ to θ^f_1 and decreasing over the range θ^c_2 to $\theta_{\max}$ then $\frac{\partial^2 E\left[C(a)\right]}{\left(\partial y_{\max}\right)^2} \geq 0$. Thus, (A7) shows us that in the symmetric collar case, simultaneously and equally increasing the reserve and buyback limits decreases costs, but at a decreasing rate as we increase $y_{\max}$ (Corollary 2).

References

Adda, J. and Cooper, R. *Dynamic Economics*: *Quantitative Methods and Applications*. 2003, MIT Press, Cambridge, MA.

Aiyagari, R.S. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 1994, *109*(3), pp. 659-84.

Burtraw, D.; Palmer, K.L. and Kahn, D.B. "A Symmetric Safety Valve." *Energy Policy*, 2010, *38*(9): 4921-32.

Deaton, A. "Saving and Liquidity Constraints." *Econometrica*, 1991, *59*(5), pp. 1221-48.

EIA. "Energy Market and Economic Impacts of H.R. 2454, the American Clean Energy and Security Act of 2009," Washington, DC: Energy Information Administration, 2009.

Fell, H. and Morgenstern, R.D. "Alternative Approaches to Cost Containment in a Cap-and-Trade System," *Resources for the Future DP 09-14* Washington, DC: Resources for the Future, 2009.

Grüll, G. and Taschini, L. "Cap-and-Trade Properties under Different Hybrid Scheme Designs." *Journal of Environmental Economics and Management*, 2011, *61*(1), pp. 107-18.

Kling, C.L. and Rubin, J. "Bankable Permits for the Control of Environmental Pollution." *Journal of Public Economics*, 1997, *64*(1), pp. 101-15.

Kopp, R., Morgenstern, R.D., Pizer, W.A. Something for Everyone: A Climate Policy That Both Environmentalists and Industry Can Live With. Weathervane September 29, 1997. Resources for the Future. Washington, D.C.

Leiby, P. and Rubin, J. "Intertemporal Permit Trading for the Control of Greenhouse Gas Emissions." *Environmental & Resource Economics*, 2001, *19*(3), pp. 229-56.

McKibbin,W.J., Wilcoxen, P.J. "A Better Way to Slow Global Climate Change." Brookings Policy Brief No. 17. Brookings Institution, Washington, DC. 1997.

Montgomery, D.W. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory*, 1972, *5*(3), pp. 395-418.

Murray, B.; Newell, R.G. and Pizer, W. A. "Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade." *Review of Environmental Economics and Policy*, 2009, *3*(1), pp. 84-103.

Newell, R.G.; Pizer, W.A. and Zhang, J. "Managing Permit Markets to Stabilize Prices." *Environmental & Resource Economics*, 2005, *31*(2), pp. 133-57.

Philibert, C. "Price Caps and Price Floors in Climate Policy," *IEA Information Paper.* Paris, France: International Energy Agency, 2008.

Pizer, W. A. "Combining Price and Quantity Controls to Mitigate Global Climate Change." *Journal of Public Economics*, 2002, *85*(3), pp. 409-34.

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

Rubin, J. "A Model of Intertemporal Emission Trading, Banking and Borrowing." *Journal of Environmental Economics and Management*, 1996, *31*(3), pp. 269-86.

Salant, S.W. "The Vulnerability of Price Stabilization Schemes to Speculative Attack." *Journal of Political Economy*, 1983, *97*(1), pp. 1-38.

Schennach, S.M. "The Economics of Pollution Permit Banking in the Context of Title." *Journal of Environmental Economics and Management*, 2000, *40*(3), pp. 189-210.

Stocking, A. "Unintended Consequences of Price Controls: An Application to Allowance Markets." Congressional Budget Office, Working Paper Series. Washington, D.C., 2010.

Weber, T.A. and Neuhoff, K. "Carbon Markets and Technological Innovation." *Journal of Environmental Economics and Management*, 2010, *60*(2), pp. 115-32.

Webster, M., Wing, I.S., and Jakobovits, L. "Second-Best Instruments for near-Term Climate Policy: Intensity Targets Vs. The Safety Valve." *Journal of Environmental Economics and Management*, 2010, *59*(3), pp. 250-59.

Weitzman, M. "Optimal Rewards for Economic Regulation." *American Economic Review*, 1978, *68*(4), pp. 683-91.

Zeldes, S.P. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics*, 1989, *104*(2), pp. 275-98.

Figures and Tables

Figure 1. Elements of the Expected Cost Function with Price Collar

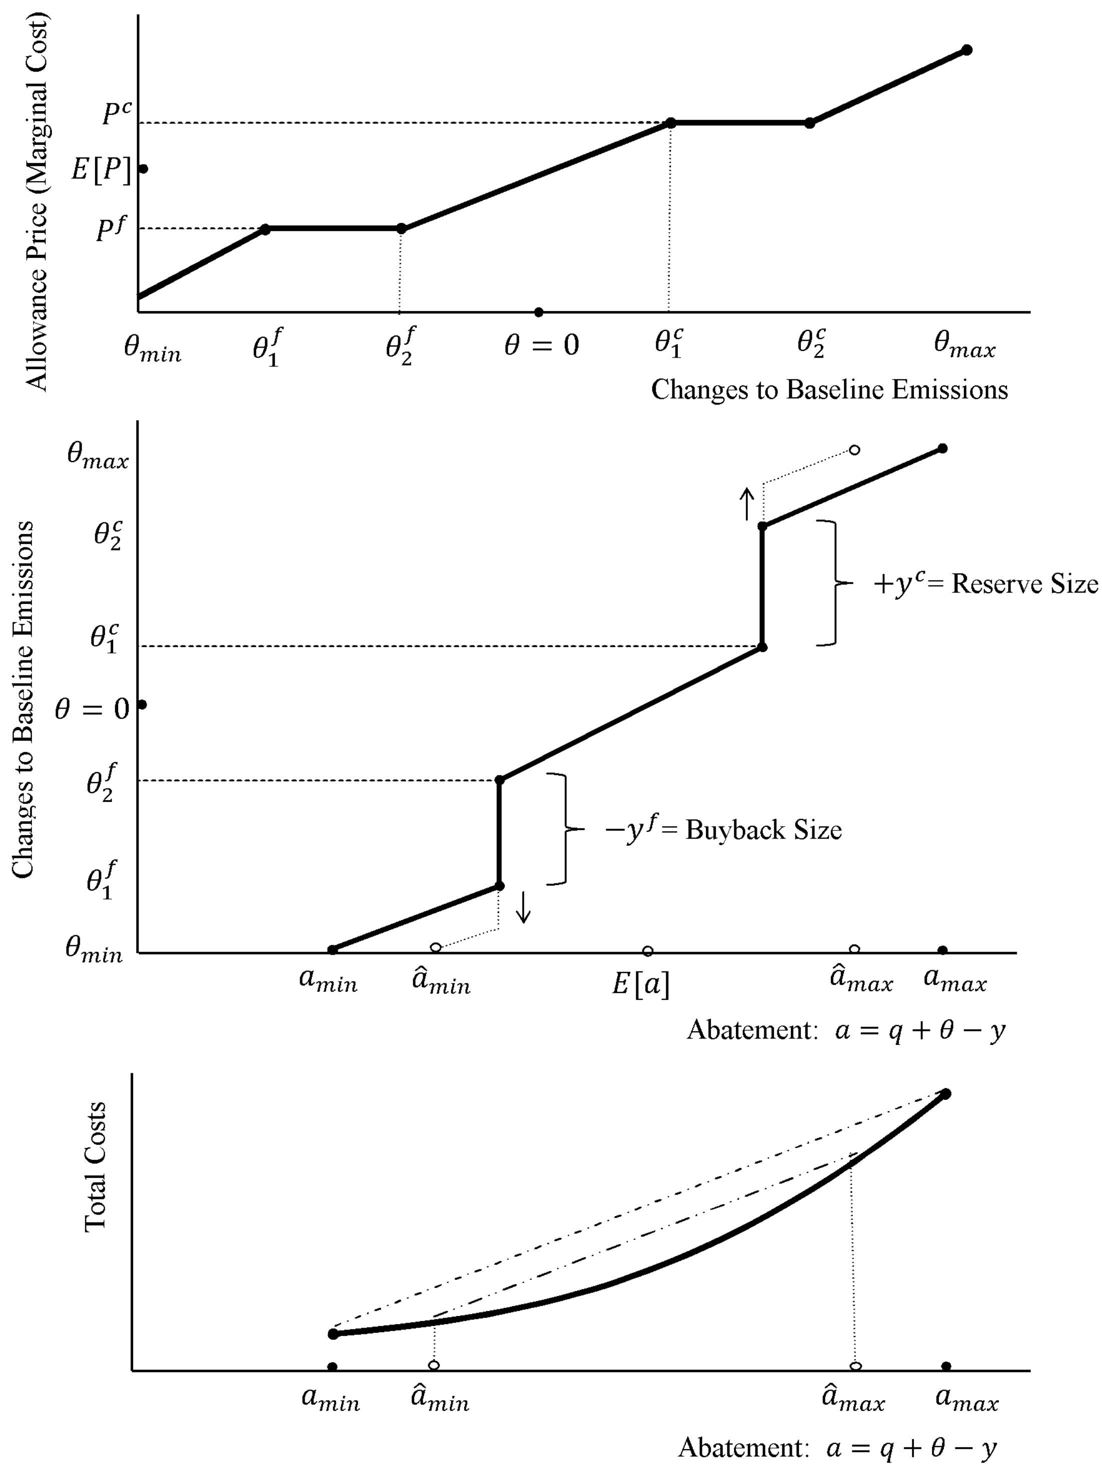

Figure 2. Share of Total Cost Savings with Soft Collar

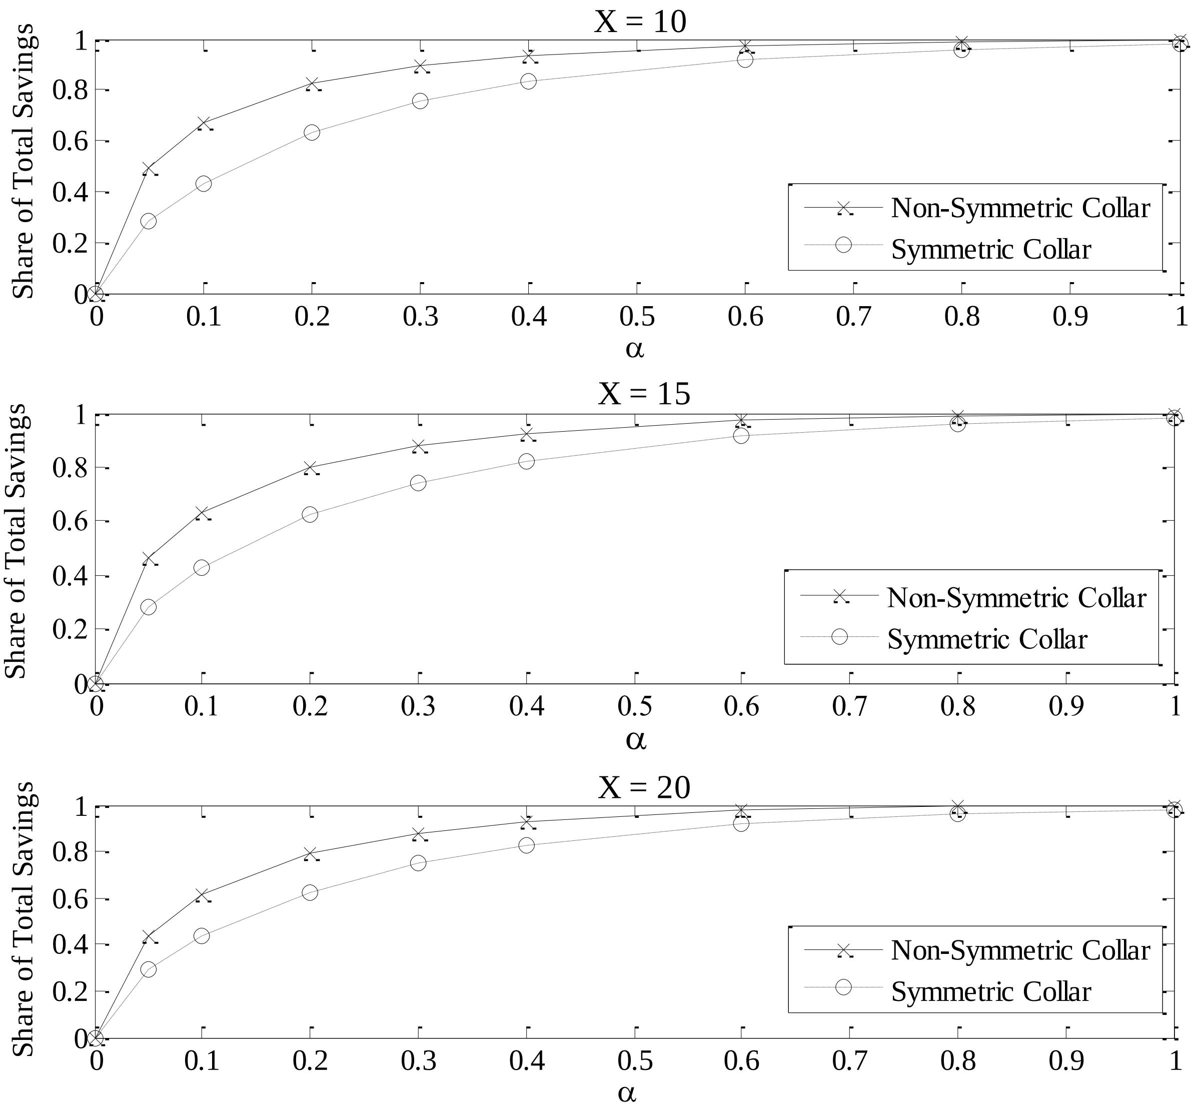

Note: This shows the share of the hard collar cost savings achieved at a particular value of α, for each value of X.

Figure 3. 95% Coverage Intervals for Emissions

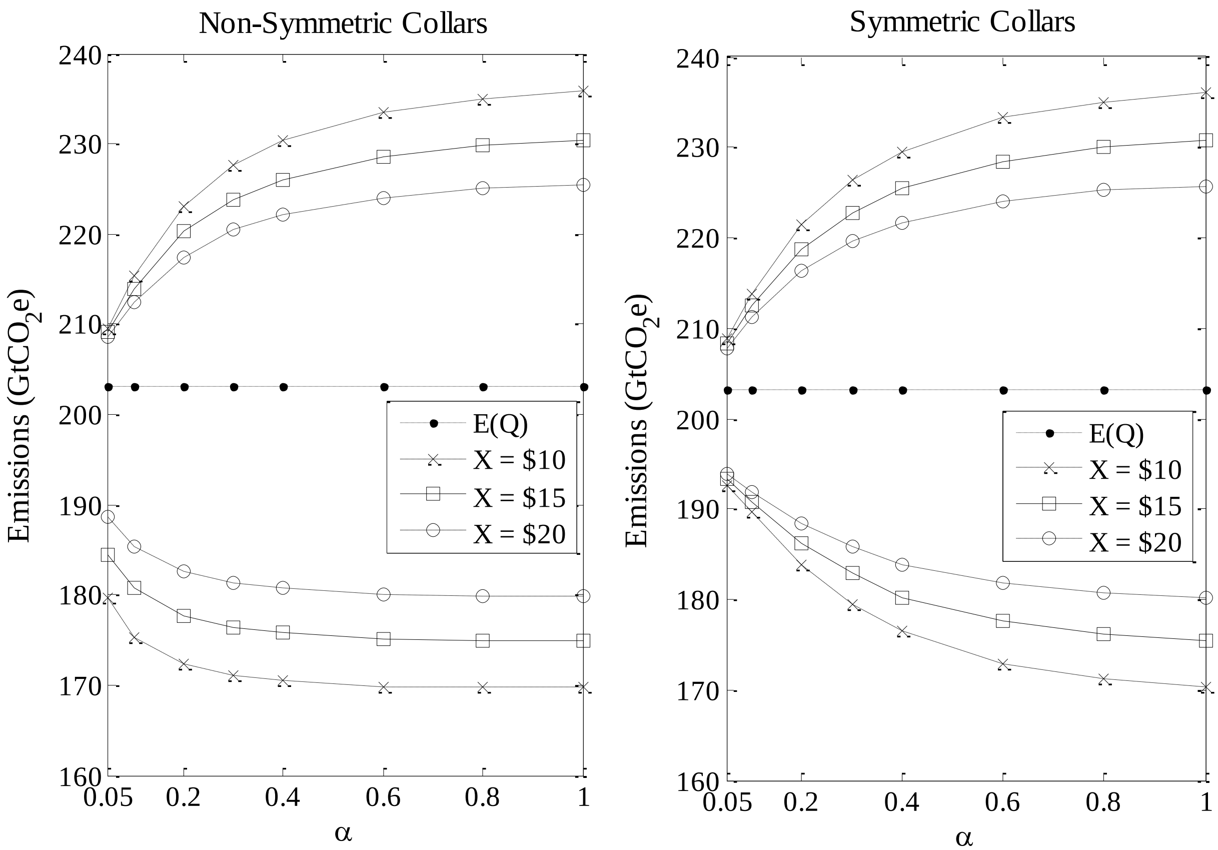

Notes: E(Q) refers to the expected cumulative emissions. All collars are adjusted so that E(Q) = 203.2 $GtCO_2e$. Because the firm will "burn" excess allowances rather than increase emissions beyond $\overline{q}_t$, some emissions uncertainty exists even when no collar is present to add or subtract allowances from the system. Hence, for the no collar case ($\alpha = 0$), E(Q) = 202.7 $GtCO_2e$, with a 95% coverage interval of [196,203].

Figure 4. Emissions Histograms for Two Soft Collars

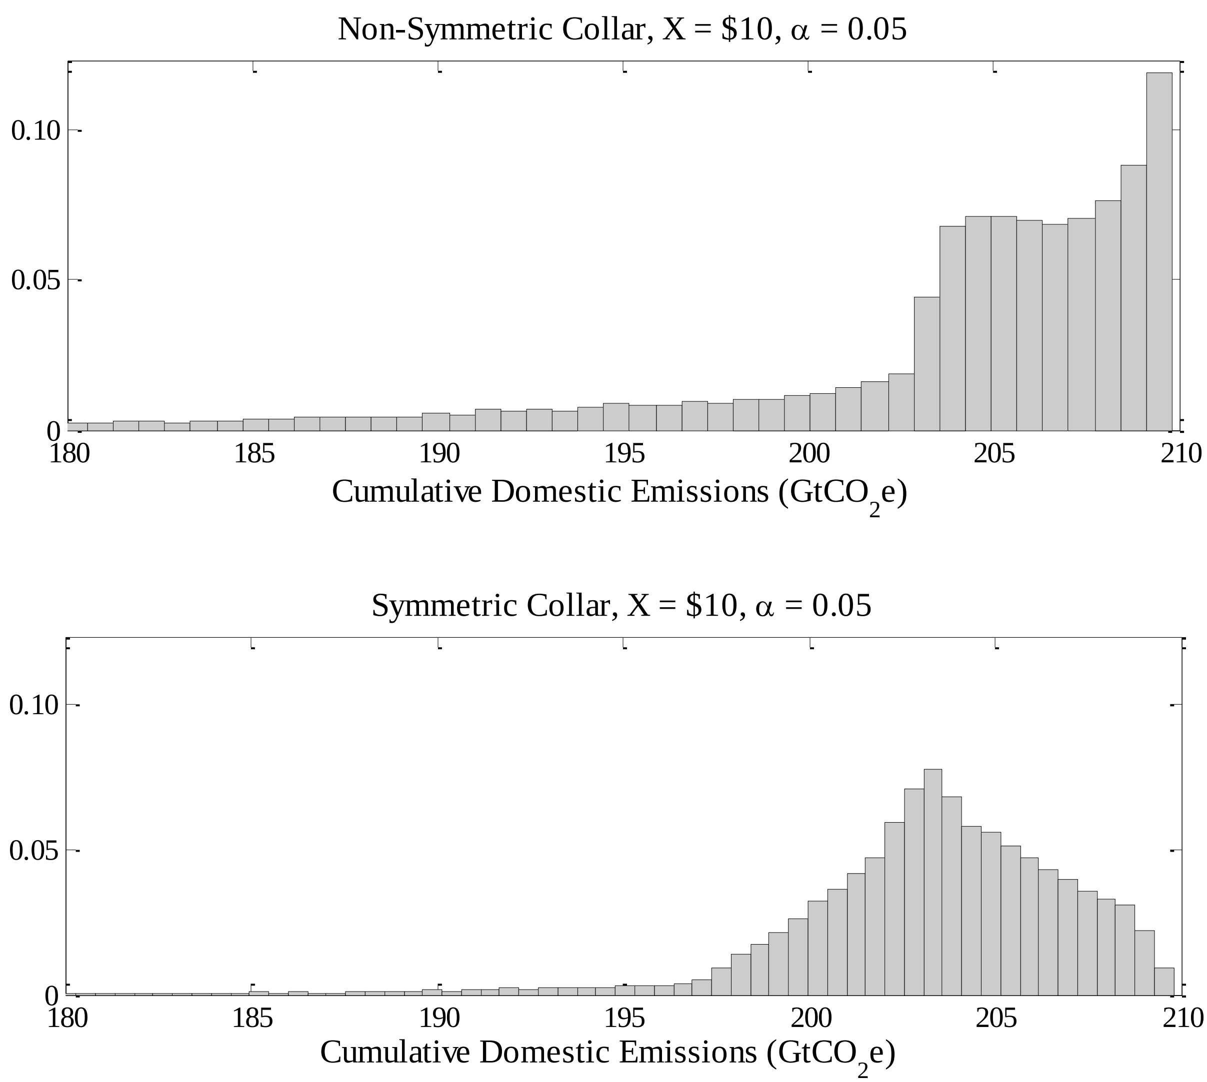

Note: This histogram give cumulative emissions over 30,000 model simulations.

Table 1: Parameter Values and Descriptions

Parameter	Value	Definition
T	39	Terminal period
β	0.952	Discount factor
ρ	0.9	AR(1) parameter for baseline emissions shock
σ_1^2	0.2	Variance of error term in baseline emissions shock
g_c	-0.0125	Rate of decline for slope of marginal abatement cost curve
c_0	\$62/tCO_2e per GtCO_2e	Initial slope of marginal abatement cost curve
$B_{min,t}$	0	Minimum bank level
y_t	-	Allowance allocation
$\overline{q}_t$	-	Expected baseline emissions
X	(\$10, \$15, \$20)	Initial width of price collar
α	(0, 0.05, 0.1, 0.2, 0.3, 0.4, ∞)	Size of reserve, as a fraction of y_t

Notes: tCO_2e stands for metric tons of CO_2 equivalent and GtCO_2e stands for giga-metric tons of CO_2 equivalent. $\alpha = 0$ and $\alpha = \infty$ represent the cases with no collar and a hard collar, respectively.

Table 2. Expected NPV of Abatement Cost

No Collar:	580.7 [46,1802]					

Collar Type: Non-symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	519.9 [133,1584]	498.1 [178,1409]	478.7 [216,1182]	470.1 [233,1035]	465.3 [242,941]	456.9 [252,694]
$15	531.7 [93,1603]	513.9 [123,1454]	496.4 [152,1257]	487.9 [166,1122]	483.1 [173,1036]	475.3 [183,840]
$20	543.1 [67,1625]	528.3 [85,1495]	513.1 [107,1315]	505.4 [117,1197]	501.1 [123,1122]	494.9 [130,989]

Collar Type: Symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	544.2 [68,1615]	524.3 [90,1458]	498.5 [125,1227]	483.5 [152,1063]	474.3 [175,961]	456.9 [252,694]
$15	549.2 [60,1633]	532.8 [77,1493]	511.1 [102,1288]	498.1 [122,1144]	489.9 [138,1050]	475.3 [183,840]
$20	553.8 [54,1644]	540.9 [65,1527]	523.2 [84,1339]	512.6 [97,1217]	505.9 [108,1132]	494.9 [130,989]

Notes: Costs given in billions of real U.S. dollars (2005 base year). Brackets indicate 95% coverage intervals from simulations. Values in the "*X*" column are the initial width of the price collar. Bold headers indicate different values of α, with $\alpha = \infty$ representing the hard collar case.

Table 3. Initial Price Ceiling levels $\left(P_1^c\right)$

Collar Type: Non-symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	19.17	21.10	22.49	23.07	23.36	23.72
$15	21.80	23.58	25.02	25.63	25.94	26.34
$20	24.46	26.23	27.69	28.30	28.62	29.02

Collar Type: Symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	23.11	23.59	23.74	23.73	23.70	23.72
$15	25.30	26.06	26.33	26.34	26.31	26.34
$20	27.28	28.47	28.94	29.00	28.99	29.02

Note: Allowance price given in $/t$CO_2$. Above values represent the P_1^c that results in expected cumulative domestic emissions of 203.2 GtCO_2e, for each parameter setting. Values in the "*X*" column are the initial width of the price collar. Bold headers indicate different values of α, with $\alpha = \infty$ representing the hard collar case.

Table 4. Root Mean Squared Error of Allowance Price Growth Rates

No Collar: 39.02

Collar Type: Non-symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	12.02	11.22	10.94	10.89	10.82	10.52
$15	14.20	13.51	12.87	12.58	12.37	11.90
$20	16.03	15.12	14.31	13.91	13.68	13.24

Collar Type: Symmetric

X	**0.05**	**0.10**	**0.20**	**0.30**	**0.40**	∞
$10	22.96	16.94	14.23	13.21	12.63	10.52
$15	24.20	18.16	15.58	14.65	14.11	11.90
$20	24.79	19.16	16.63	15.77	15.26	13.24

Notes: Allowance price given in $/t$CO_2$. Values in the "*X*" column are the initial width of the price collar. Bold headers indicate different values of α, with $\alpha = \infty$ representing the hard collar case.